

Auction Mechanisms for Distributed Spectrum Sharing*

Jianwei Huang, Randall A. Berry, Michael L. Honig

Department of Electrical and Computer Engineering

Northwestern University

Evanston, IL 60208

{jianweih,rberry,mh}@ece.northwestern.edu

Abstract

We study auction mechanisms for allocating power among a group of spread spectrum users. The users are assumed to share the bandwidth with a licensed user, or spectrum owner, which imposes a received power constraint (corresponding to a constraint on interference) at a particular measurement location. Both co-located and non-located receivers are considered. Each user receives a utility that is a function of the received Signal-to-Interference plus Noise Ratio (SINR). We propose two auction mechanisms for power allocation in which the spectrum owner charges for SINR and received power, respectively, and compare the associated utility and revenue achieved in some simple cases. We also derive an iterative and distributed bid updating algorithm, and specify conditions for which this algorithm converges globally to the Nash Equilibrium (NE) of the auction.

1 Introduction

Recent advances in software-defined radio have made it possible to share both licensed and unlicensed radio spectrum more efficiently. The current “command-and-control” policy for spectrum allocation does not exploit this capability, and is stimulating investigation into more flexible market-based mechanisms for dynamic sharing spectrum.

In this paper, we study spectrum sharing subject to an *interference temperature* constraint, which is a constraint on the RF power per unit bandwidth measured at a particular location [1]. This is motivated by the scenario in which users (or service providers) wish to share spectrum owned by, or licensed to another primary user, such as another service provider or government agency. The primary user determines the interference temperature constraint(s) so as not to disrupt its own service. We assume that the primary user provides a spectrum *manager* to control the amount of bandwidth and power assigned to users sharing the spectrum. The users may transmit to receivers at different locations, or to co-located receivers at a single access point.

This paper is an extension of prior work [2] in which auction mechanisms are studied for allocating power among spread spectrum users. Namely, each user’s transmitted power is uniformly spread across the entire available bandwidth controlled by the manager, so that the combined power-bandwidth allocation problem reduces to a received power

*This work was supported by the Northwestern-Motorola Center for Communications, and by NSF CAREER award CCR-0238382.

allocation problem. Each user has a utility function, which depends on the received Signal-to-Interference plus Noise Ratio (SINR), reflecting the desired Quality of Service (QoS).

As in [2], we consider simple auctions mechanisms for power allocation. (See also [3,4], which consider similar mechanisms for wire-line resource allocation.) Namely, the manager announces a price (for either received power or SINR), the users submit bids for the amount of power they wish to purchase, and the manager allocates power proportional to the bids received. In contrast to [2], which considers a multiple-access scenario in which the measurement point and receivers are co-located, here we allow the measurement point and receivers to be at different locations. This includes a peer-to-peer scenario in which the different receivers are not co-located. The interference a user receives therefore depends on the other users' transmit powers, cross-channel gains, as well as the bandwidth. We also extend some of the results in [2] for co-located receivers. For example, we compare the revenue generated for pricing of power versus pricing of SINR.

We first analyze these auctions as a simultaneous move game [5], assuming all information (i.e., utility functions and link channel gains) is available to the users, but not to the manager. We give conditions under which this game has a unique Nash Equilibrium (NE) and characterize this NE. We subsequently formulate an iterative and distributed algorithm for power allocation, in which each user needs to know only the channel gains associated with the links connected to that user. We show that this algorithm converges globally to the unique NE of the simultaneous move game, when it exists. The auction mechanisms considered are therefore scalable with the population size.

In the next section we present the system model and the auction mechanisms considered. We start with the Vickrey-Clarke-Groves (VCG) auction, which achieves a socially optimal allocation, i.e., maximizes total utility [5]. To reduce the amount of information exchange required and computational complexity, we present simpler *share* or *divisible* auction mechanisms [4]. We then state our results on the properties of the NE of these auctions, which generalize the results in [2] to the general case in which receivers are not co-located. We evaluate the revenue collected by the manager in Sec. 5, assuming co-located receivers. In Sec. 6 we give an iterative and distributed bid updating algorithm, and characterize its convergence. Illustrative numerical results are given in Sec. 7 and conclusions are presented in Sec. 8.

2 Auction Mechanisms

2.1 System Model

A block of spectrum with bandwidth B is to be shared among M spread spectrum users. In what follows, we will consider the two scenarios shown in Fig. 1. Namely, Fig. 1 (a) shows co-located receivers, and Fig. 1 (b) shows non-co-located receivers. The users share the spectrum with a primary user, who imposes a constraint on the total interference received at a measurement point (node M in Fig. 1). In general, the measurement point may or may not be co-located with a receiver.

It will be convenient to denote each transmitter-receiver *pair* by a single index i , which we will refer to as a particular "user". Assuming each transmitter spreads power uniformly across the band, the received SINR for user i is given by

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

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

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

where the index 0 corresponds to the measurement point.

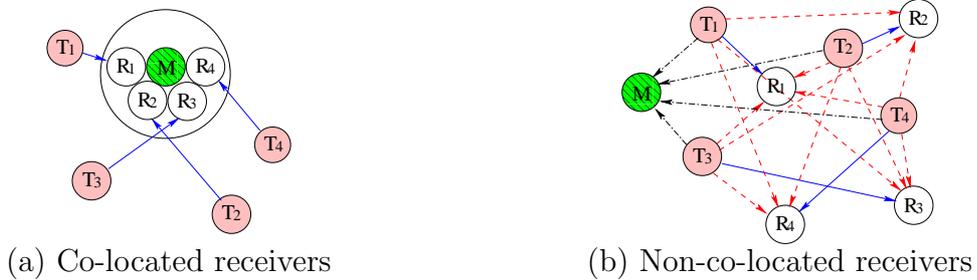


Figure 1: Network scenarios: node M is the measurement point, and nodes T_i and R_i are, respectively, the transmitter and receiver associated with user i .

Each user i is assigned a utility function $U_i(\gamma_i) = U_i(\theta_i; \gamma_i)$, which reflects the utility derived by the user as a function of the received SINR. Here θ_i is a user-dependent priority parameter. As an example, the *logarithmic utility function* $U_i(\gamma_i) = \theta_i \ln(\gamma_i)$.¹ We will assume that 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 [6].

A power allocation is *socially optimal* if it maximizes the total (sum) utility $\sum_{i=1}^M U_i(\gamma_i)$. A socially optimal allocation is *Pareto optimal*, i.e., no user’s utility can be increased without decreasing another user’s utility. However, Pareto optimality does not imply social optimality. We also note that a power allocation is Pareto optimal if and only if the total received power constraint is tight, i.e., $\sum_{i=1}^M p_i h_{i0} = P$, hence this is a necessary condition for social optimality.

We assume that each user’s utility function is private information, i.e., is initially known only to that user. 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. In auction theory, a well known mechanism for achieving a socially optimal allocation is the generalized VCG auction.

2.2 VCG Auction

In our context, a VCG auction can be described as follows: First, users are asked to submit their utility functions $\{U_i(\gamma_i)\}$.² The manager then computes the maximum total utility $U_{\max} = \max_{\{p_j\}} \sum_{j=1}^M U_j(\gamma_j)$ 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 $i \in \{1, \dots, M\}$. In total, the manager must solve $M + 1$ optimization problems. The

¹This approximates the weighted rate of user i in the high SINR regime.

²The users are not obligated to submit the correct utility functions, i.e., they may lie about their utility.

manager then charges user i the amount $U_{\max} - U_{\max/i}$, which is the incremental social benefit derived from including user i in the auction.

It can be shown that this mechanism results in a socially optimal outcome, and it is a (weakly) dominant strategy for users to bid truthfully (i.e., state their true utility functions). The VCG auction may not be suitable for dynamic spectrum sharing for the following reasons: (i) To specify the users' utility functions, in particular, the SINR in (1) for each user i , the channel gains h_{ij} for all $i, j \in \{1, \dots, M\}$ must be measured by the users and reported to the manager. (ii) The manager must solve $M + 1$ optimization problems, which are typically non-convex due to the interference. Hence the information exchange and computational requirements are likely to become excessive for large M . We therefore examine simpler mechanisms for power allocation.

2.3 One-Dimensional Share Auctions

We now describe a share, or divisible auction for power allocation in which the users submit one-dimensional bids for the amount of resource they wish to purchase at an announced price, and the manager simply allocates the received power in proportion to the bids. Each user then pays the announced price times the amount of allocated resource (i.e., power or SINR). We also assume that the manager announces a reserve bid $\beta \geq 0$, and transmits with the corresponding reserve power, which interferes with the other users. Here the main purpose of the reserve bid is to guarantee a unique desirable outcome of the auction, rather than to extract more revenue from the other bidders [7]. We will show that the interference generated by the manager can be made arbitrarily small.

We first present an auction mechanism, which assumes complete information, i.e., all users' utility functions and all channel gains are known to all users. In Sec. 6, we present a distributed bid updating algorithm that only requires limited information, i.e., each user i needs to know only the channel gains $\hat{h}_{ii} = h_{ii}/h_{i0}$ 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 $i \in \{1, \dots, M\}$ submits a bid $b_i \geq 0$.
3. The manager keeps reserve power p_0 , and allocates to each user i a transmit power p_i so that the received power at the measurement point is proportional to the bids, i.e.,

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

The received SINR for user i is given by (1) where the interference term contains $p_0 h_{0i}$, and h_{0i} is the channel gain from the manager (measurement point) to user i 's receiver. If $\sum_{i=1}^M b_i + \beta = 0$, then $p_i = 0$.

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

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

$$S_i(b_i; b_{-i}) = U_i(\gamma_i(b_i; b_{-i})) - C_i(b_i; b_{-i}). \quad (4)$$

Here we omit the dependence on β and π .

Under this mechanism the users play a simultaneous move game, where their strategies consist of the bids they submit and their pay-offs are the resulting surplus. An NE of this game is a bidding profile \mathbf{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

$$\mathcal{B}_i(b_{-i}) = \left\{ \hat{b}_i \mid \hat{b}_i = \arg \max_{b_i \in [0, \infty)} S_i(b_i; b_{-i}) \right\}, \quad (5)$$

i.e., the set of b_i 's that maximize $S_i(b_i; b_{-i})$ given a fixed b_{-i} .³ The NE bidding profile \mathbf{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 auctions for network resource allocation (e.g., [3,4]) in that the bids are not generally the same as the payments. The manager can therefore influence the NE by choosing β and π . This alleviates the typical inefficiency associated with the NE, and leads to Pareto optimal (and in some cases socially optimal) solutions.

3 SINR Auction

In this case, $C_i(\gamma_i) = \pi^s \gamma_i$ so that each user's payment depends on both the transmit power and the interference. Here we assume logarithmic utility functions, although the analysis applies to other utility functions as well (e.g., $\theta_i \log(1 + \gamma_i)$).

Theorem 1 *In an SINR-based auction with logarithmic utility functions:*

- (i) *For $\beta > 0$, there exists a threshold price $\pi_{th}^s > 0$ such that a unique NE exists if $\pi^s > \pi_{th}^s$; otherwise, no NE exists.*
- (ii) *For $\beta = 0$, there are either an infinite number of Nash Equilibria, or no NE.*

Proof. (outline): Substituting $U_i(\gamma_i) = \theta_i \log(\gamma_i)$ in (4) and setting the derivative with respect to b_i equal to zero gives

$$\mathbf{b} = \mathbf{K}\mathbf{b} + \beta \mathbf{k}_0 \quad (6)$$

where \mathbf{b} is the $M \times 1$ vector of best response bids across users, \mathbf{K} is an $M \times M$ matrix, which depends on all channel gains and the θ_i parameters, and \mathbf{k}_0 is an $M \times 1$ vector, which depends on the channel gains h_{0i} , h_{i0} and h_{ii} . The number of Nash equilibria is therefore the number of positive solutions to (6). Statement (i) in the theorem follows from an application of the Perron-Frobenius theorem [8]. If $\beta = 0$, then $(\mathbf{I} - \mathbf{K})\mathbf{b} = 0$, for which there is either no feasible solution (i.e., $\mathbf{b} = 0$), or an infinite number of solutions (i.e., if \mathbf{K} is singular). ■

Theorem 1 and the following two propositions were proven in [2] for the case of co-located receivers using different methods, which do not directly apply to the non-co-located scenario.

Note that β does not affect the power allocation at the NE since all equilibrium bids are proportional to β . The manager can therefore announce an arbitrary $\beta > 0$. This observation also follows from (3), and so applies to any utility function.

The relation (6) is similar to power control updates for CDMA, which have been analyzed, for example, in [9]. A key difference between that work and the auction model considered here is that here we consider elastic data traffic without a preset target SINR.

³In general the best response set may contain more than one element.

Also, the price can be adjusted so that a unique NE always exists. In contrast, in the CDMA model the power control updates may not converge if the target SINR is too high. The mathematical similarity arises from the fact that by designing appropriate auction mechanisms, we convert the constrained power allocation problem into an unconstrained noncooperative game, in which each user updates his bid in an attempt to reach the desired equilibrium SINR level.

An allocation $\{x_i\}_{i \in \{1, \dots, M\}}$ is *weighted max-min fair* with weights $\{w_i\}_{i \in \{1, \dots, M\}}$ if for each $i \in \{1, \dots, M\}$, x_i can not be increased without decreasing some x_j , $j \in \{1, \dots, M\}$, for which $x_j/w_j \leq x_i/w_i$.

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

We call a system *stable* if there exists a unique NE. In a stable system, define the *system usage* by $\eta = \sum_{i=1}^M p_i^* h_{i0} / P = \sum_{i=1}^M b_i^* / (\sum_{i=1}^M b_i^* + \beta)$. For Pareto optimality $\eta = 1$, but the necessary condition for stability is $\eta < 1$ since the reserve bid β must be positive. Hence Pareto optimality and stability are conflicting objectives⁴.

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

Proposition 2 *In an SINR-based auction with logarithmic utility, there exists a unique price $\pi^{s\varepsilon}$ for any $\varepsilon \in (0, 1)$, such that the system is stable and achieves an ε -Pareto optimal solution (i.e., $\eta = 1 - \varepsilon$ in the original system).*

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, the stability conditions in Theorem 1 may be violated.

4 Power Auction

Here we focus on co-located receivers, as shown in Fig. 1 (a), where $h_{ij} = h_{i0}$ for all $i, j \in \{1, \dots, M\}$. We denote user i 's received power as $p_i^r = p_i h_{i0}$. Given a Pareto optimal allocation, we have for each user i ,

$$\gamma_i \equiv \gamma_i(p_i^r) = \frac{p_i^r}{n_0 + \frac{1}{B}(P - p_i^r)}, \quad (7)$$

which, in contrast to the more general SINR expression (1), does not depend on the distribution of powers among the other users.

We say that a power allocation is ε -socially optimal if it maximizes the total utility of the ε -system. It is shown in [2] that there exists a price $\pi^{p\varepsilon}$ such that the system is stable and the NE achieves ε -social optimality for any $\varepsilon \in (0, 1)$ provided that for each $i \in \{1, \dots, M\}$, $U(\theta_i, \gamma_i)$ satisfies

$$\frac{|U_i''(\gamma_i)|}{U_i'(\gamma_i)} (\gamma_i + B) > 2, \quad (8)$$

⁴Here we do not include the power allocated to the manager in our definition of Pareto optimality.

for any $\gamma_i \in [0, P/n_0]$. Inequality (8) guarantees that the utility $U_i(\gamma_i(p_i^r))$ is strictly concave in the received power p_i^r . For many common utility functions, this condition is satisfied when the bandwidth is large enough, which implies that the interference among users is relatively small.

When (8) is not satisfied, the utility may not be concave in the received power, and the power auction may not be able to achieve an η close to 1. An example is shown in Fig. 2 with two users and logarithmic utility functions. Total demand for power (from the best response function) is plotted versus price. The maximum efficiency $\eta = 0.41$ is achieved at $\pi^p = 0.935$. At a lower price the demand exceeds the power constraint so that the auction cannot converge to an NE. In contrast, an SINR auction can achieve any η , which may lead to higher total utility.

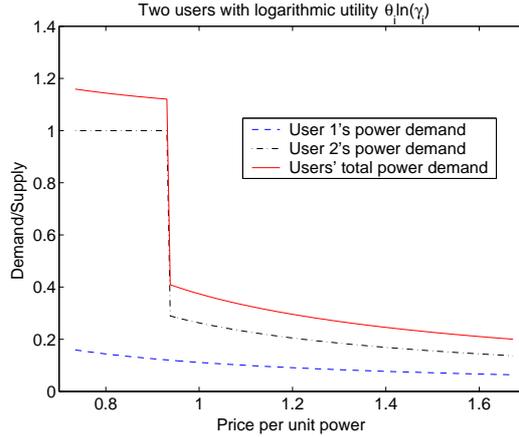


Figure 2: Demand for power versus price with logarithmic utility functions (power auction). $M = 2$, $P = 10$, $B = 1000$, $n_0 = 10^{-3}$, $\theta_1 = 1$, $\theta_2 = 2$. The condition (8) is not satisfied.

5 Revenue Comparison

Here we compare the revenue obtained from power and SINR auctions for co-located receivers.⁵ We also assume that all users have the same utility function, which satisfies (8). Let R^p and R^s be the revenue derived from power and SINR auctions, respectively.

Theorem 2 *If a power and SINR auction achieve the same usage η , then $R^p > R^s$ and $R^p/R^s \rightarrow 1$ as $M \rightarrow \infty$.*

Numerical examples indicate that Theorem 2 is still valid when the users have different utilities (see Fig. 3). If (8) is not satisfied, then the power auction may collect lower revenue than the SINR auction, since the former may not be able to achieve a particular η close to one. However, Theorem 2 remains valid for logarithmic utility functions.

Proposition 3 *Given logarithmic utility functions, assume that there exists a finite $\bar{\theta}$ such that $\theta_i < \bar{\theta}$, $1 \leq i \leq M$, for all M . Then $R^p > R^s$ and $R^p/R^s \rightarrow 1$ as $M \rightarrow \infty$.*

Notice that Proposition 3 does not assume that all users have the same utility function, or that both auctions give the same η . Hence with logarithmic utilities the power auction always generates more revenue than the SINR auction.

⁵We note that other auction mechanisms may provide more revenue.

6 Distributed Bid Updating Algorithm

In Sec. 2, we assumed that the users' utility functions and channel gains are known to the other users, so that the auction can be analyzed as a simultaneous-move game with complete information. Here we relax this assumption, and present an iterative and distributed algorithm that converges to the NE of the simultaneous auction. In this section we consider only an SINR auction with logarithmic utilities, but note that the following results also apply to other scenarios, e.g., power auction with co-located receivers. (In what follows the receivers do not have to be co-located.)

As an example, consider the SINR auction with logarithmic utilities. Suppose the users iteratively update their bids according to the best response (5), i.e.,

$$\mathbf{b}^{(t)} = \mathbf{K}\mathbf{b}^{(t-1)} + \mathbf{k}_0\beta, \quad (9)$$

for $t = 1, 2, \dots$, where $\mathbf{b}^{(0)}$ is an arbitrary (positive) initial bidding profile. Clearly, if the bid profile converges, it will converge to the unique solution to (6), if it exists. Furthermore, it can be shown that when (6) has a unique solution, then \mathbf{K} will have a spectral radius less than 1 and so (9) will converge from any initial bid profile.

Proposition 4 *For the SINR auction with logarithmic utility functions, (9) is equivalent to the set of individual user updates*

$$b_i^{(t)} = \frac{\theta_i/\pi^s - \gamma_i^{(t-1)}\varphi_i}{\gamma_i^{(t-1)} - \gamma_i^{(t-1)}\varphi_i} b_i^{(t-1)}, \quad (10)$$

for $1 \leq i \leq M$, where $\varphi_i = n_0\theta_i / (\hat{h}_{ii}P\pi^s)$.

The update (10) requires that user i know only the channel gain ratio h_{ii}/h_{i0} , the background noise n_0 , and the received SINR $\gamma_i^{(t)}$ at each iteration. Hence this mechanism is distributed and scalable, and, from the above discussion, it will converge to the unique NE of the simultaneous auction, if it exists. In practice, the manager would adaptively search for a price π^s a little higher than π_{ih}^s (unknown *a priori*), which guarantees a unique NE, and achieves an efficiency η close to one.

7 Numerical Results

Fig. 3 compares utility and revenue for SINR and power auctions for co-located receivers. The total power constraint $P = 10$, the background noise power $n_0 = 10^{-3}$, and users have the rate utility $\theta_i \log(1 + \gamma_i)$, where θ_i is a uniformly distributed random variable in $[1, 100]$. Fig. 3(a) shows the ratio of revenue for the power auction to the revenue for the SINR auction versus $Bn_0/(P + 2n_0)$, which is the left side of the condition (8), and indicates the degree to which $U_i(\cdot)$ is a non-concave function of the transmit powers. Fig. 3(b) shows the corresponding ratio of utilities. Each point is averaged over 100 independent realizations. When condition (8) is satisfied, corresponding to $Bn_0/(P + 2n_0) > 0$ dB, the target efficiency $\eta^* = 1 - 10^{-6}$ can be achieved. In that case, the power auction generates more revenue than the SINR auction, and the ratio diminishes to one as M increases. When (8) is not satisfied ($Bn_0/(P + 2n_0) < 0$ dB), the manager adjusts the price to achieve the highest possible η , and the revenue ratio could be less than one.

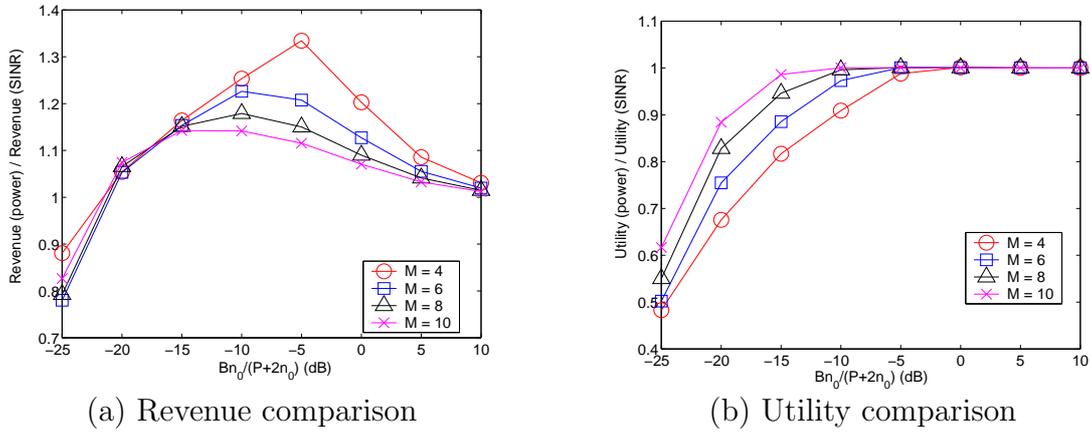


Figure 3: Comparison of revenue and utility for power and SINR auctions with rate utility functions and co-located receivers.

Fig. 3(b) shows that the utility ratio is essentially one when (8) is satisfied, since both auctions achieve close to the maximum total utility. When $Bn_0/(P + 2n_0) < 0$ dB the utility ratio is less than one, since the power auction cannot achieve an η close to one.

Figs. 4 (b) and (c) show the convergence of users' bids and transmit powers in an SINR auction using the distributed algorithm in Sec. 6 for the network shown in Fig. 4(a). The network has three users and non-co-located 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 1 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. Fig. 4(c) shows that after compensating for the interference, user 2 transmits with the highest power, and user 1 transmits with the lowest power.

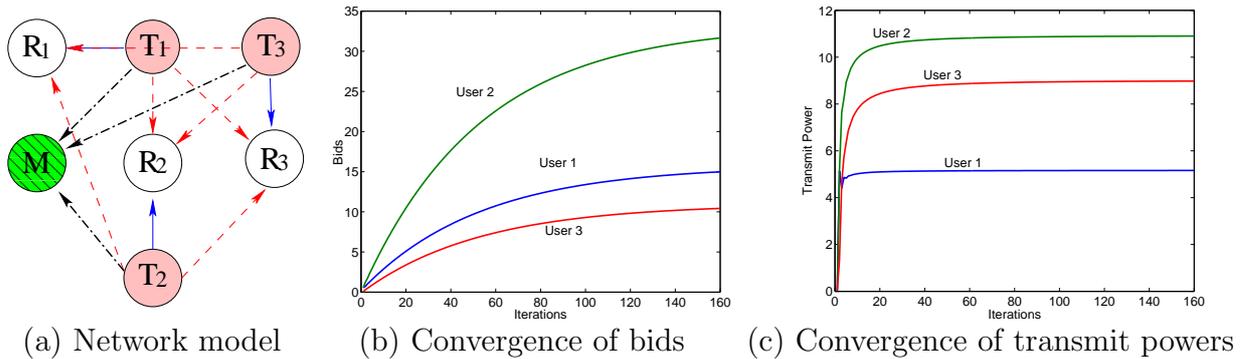


Figure 4: Transient behavior of SINR auction in a three-user network with non-co-located receivers and logarithmic utility function.

8 Conclusions

We have studied auction mechanisms for distributed spectrum sharing among a group of spread spectrum users with non-co-located receivers. The interference temperature, or total received power, is constrained at a particular measurement point. For the SINR auction with logarithmic utility functions, a large enough price guarantees a unique NE (provided that the manager announces a positive reserve bid). With co-located receivers, general utility functions, and large enough bandwidth, the power auction achieves the socially optimal allocation. (The SINR auction was observed numerically to give similar performance.) This is because a large bandwidth implies that the utilities are concave functions of the transmit powers. Also, in this scenario the power auction generates more revenue than the SINR auction, although the ratio diminishes to one as the number of users becomes large. The utilities may not be concave for smaller bandwidths, in which case the power auction cannot achieve usages η close to one. In that case, we observed numerically that the total utility and revenue for the power auction can be lower than the utility and revenue for the SINR auction. We also presented an iterative, distributed bid updating algorithm, which requires that each user know only local information, and converges globally to the NE.

References

- [1] "Spectrum policy task force report," Federal Communications Commission, US, Nov. 2002.
- [2] J. Huang, R. Berry, and M. L. Honig, "Auction-based spectrum sharing," in *the 2nd Workshop on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, University of Cambridge, UK, March 2004.
- [3] 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.
- [4] R. Maheswaran and T. Basar, "Nash equilibrium and decentralized negotiation in auctioning divisible resources," 2003, Group Decision and Negotiation, to appear.
- [5] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford University Press, 1995.
- [6] S. Shenker, "Fundamental design issues for the future internet," *IEEE Journal of Selected Areas in Communications*, vol. 13, pp. 1176–1188, 1995.
- [7] P. Milgrom, *Putting Auction Theory to Work*. Cambridge University Press, 2004.
- [8] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.
- [9] 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.