

Strategyproof Auctions for Balancing Social Welfare and Fairness in Secondary Spectrum Markets

Ajay Gopinathan, Zongpeng Li

Department of Computer Science, University of Calgary
 {ajay.gopinathan, zongpeng}@ucalgary.ca

Chuan Wu

Department of Computer Science, University of Hong Kong
 cwu@cs.hku.hk

Abstract—Secondary spectrum access is emerging as a promising approach for mitigating the spectrum scarcity in wireless networks. Coordinated spectrum access for secondary users can be achieved using periodic spectrum auctions. Recent studies on such auction design mostly neglect the repeating nature of such auctions, and focus on greedily maximizing social welfare. Such auctions can cause subsets of users to experience starvation in the long run, reducing their incentive to continue participating in the auction. It is desirable to increase the diversity of users allocated spectrum in each auction round, so that a trade-off between social welfare and fairness is maintained. We study truthful mechanisms towards this objective, for both *local* and *global* fairness criteria. For local fairness, we introduce randomization into the auction design, such that each user is guaranteed a minimum probability of being assigned spectrum. Computing an optimal, interference-free spectrum allocation is NP-Hard; we present an approximate solution, and tailor a payment scheme to guarantee truthful bidding is a dominant strategy for all secondary users. For global fairness, we adopt the classic max-min fairness criterion. We tailor another auction by applying linear programming techniques for striking the balance between social welfare and max-min fairness, and for finding feasible channel allocations. In particular, a pair of primal and dual linear programs are utilized to guide the probabilistic selection of feasible allocations towards a desired tradeoff in expectation.

I. INTRODUCTION

Networking applications that rely on wireless technology have enjoyed rapid growth in recent years. The proliferation of wireless devices naturally results in an increase in demand for usable spectrum, a scarce, government controlled commodity. Spectrum availability and usage are usually regulated by governmental agencies such as the Federal Communications Commission (FCC) in the United States. Traditionally, spectrum allocation has been conducted through the use of large scale auctions [1]. Recent studies show that such a *static* allocation of spectrum is inefficient, with spectrum utilization varying dramatically in both location and time [2], [3]. Consequently, spectrum sharing through the use of a dynamic, real-time secondary spectrum market has been proposed for mitigating the problem of spectrum scarcity [4], [5]. In such a market, primary users periodically hold *auctions* in order to lease idle portions of its spectrum to unlicensed secondary users. In fact, such markets are no longer merely theoretical, as proved by the advent of companies like Spectrum Bridge [6], who operate online marketplaces for selling spectrum leases, through auctions along with other means.

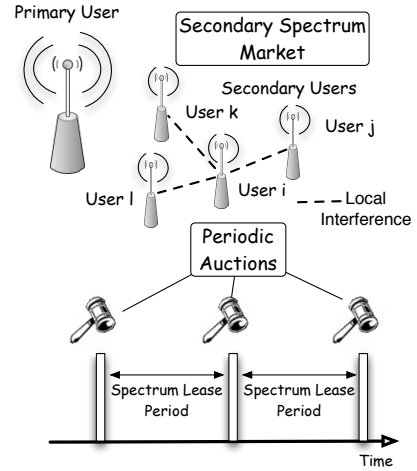


Fig. 1. A secondary spectrum market – the primary user holds auctions periodically to lease idle portions of spectrum to secondary users.

Fig. 1 shows a *secondary spectrum market*, which consist of a primary user, who owns the spectrum license, and a number of unlicensed secondary users. The spectrum utilization for the primary user varies dynamically, and at any given time, the primary user may have idle portions of spectrum. Due to spectrum demand from secondary users, the primary user pools this unused spectrum, and leases out chunks (channels) to secondary users for short periods of time, through auctions. A natural goal of such spectrum allocation is maximizing the *social welfare*, *i.e.*, allocating spectrum to users who value it the most. Unlike traditional auctions, secondary spectrum auctions are characterized by two unique properties. First, a single chunk of spectrum, or a channel, has the possibility for spatial reuse, and may be leased to multiple secondary users, as long as they do not *interfere*. Second, the temporal dynamics of spectrum usage by the primary user implies that these leases are necessarily ephemeral [4], [7]–[9], and hence the auction must be repeated periodically. Both properties lead to interesting new challenges for the design of secondary spectrum auctions.

Given that spectrum auctions are repeated, achieving some form of fairness along the temporal dimension is desired, when computing the spectrum allocation. An auction that greedily maximizes the total utility of all users (*i.e.*, the *social welfare*) in each round could lead to a subset of secondary users

starving for spectrum. Starvation is not only unsatisfactory from the perspective of *fairness*, but also has other undesirable consequences in repeated auctions. First, it discourages losing users from continuing to participate in the auction [10], which is in general detrimental to the revenue for the primary user and to social welfare in the long run. Second, bidder starvation can lead to *vindictive bidding* [10], [11], in which users with no chance of winning increase their bids, causing winning users to pay a higher price. In fact, such behaviour has been observed in as much as 40% of bidders in Yahoo’s online ad auctions [11]. Such harmful effects can be mitigated by *increasing the diversity* of the set of users who win spectrum, in each auction round. In conclusion, social welfare maximization in spectrum auctions should be curbed by fairness constraints.

Let us make the argument for fairness in a spectrum allocation setting more concrete. Consider the simple topology of four interfering nodes, as shown in Fig. 1. Assume that all the agents i, j, k and l , each has the valuation of V dollars for the channel, compete through an auction. If the auction greedily maximizes social welfare in each round, it always allocates spectrum to every agent except for i . In fact, the only way for i to be guaranteed spectrum is to bid higher than $3V$, which is 3 times as much as other agents’ valuation in the system. In this scenario, it’s not irrational for i to simply drop out of the auction, which may reduce the primary user’s revenue to 0, *e.g.*, under Vickrey pricing [12]. Observe that due to interference, i is forced to compete with the *combined valuation* of agents j, k and l , whereas each of the latter competes with i alone. In a repeating auction, we can mitigate this form of unfairness and provide motivation for i to continue bidding in each auction round, essentially by letting i win the auction occasionally.

For the reasons stated above, it’s desirable to increase the diversity of the winning set of users, so that there is a trade-off between maximizing social welfare and providing a minimum service guarantee. Designing an auction that achieves this goal, without compromising the property of incentive-compatibility (*i.e.* truthfulness) is precisely the goal of our work. We design auction mechanisms that compute an interference-free spectrum allocation that maximizes social welfare, subject to fairness constraints. Our mechanisms are flexible in that it provides the primary user with the freedom to choose a trade-off between social welfare and fairness. In repeating secondary spectrum auctions, our mechanisms provably achieve the required trade-off *in expectation*. In particular, we provide two mechanisms for achieving this goal, targeting local and global fairness, respectively. The first mechanism is a truthful, *randomized* auction framework that can be used to provide fairness in the form of either a fixed minimum probability of allocation for each agent, or a restricted notion of envy-free fairness [13], in which users are allocated spectrum if their valuation matches that of their neighbours in the conflict graph. Since computing an interference-free spectrum allocation is NP-Hard, we resort to an approximation technique. It is known that the Vickrey-Clarke-Groves (VCG) [12], [14], [15] payment scheme is

not truthful, when applied to approximate solutions [4], [16]. Consequently, we tailor an alternate payment scheme to ensure users have no incentive to lie when bidding.

The previously described mechanism is useful for implementing fairness measures that are “local”, in the sense that we provide a minimum probability of spectrum allocation by prioritising users with respect to its neighbours. In the event that the primary user is interested in providing a more global measure of fairness, such as max-min fairness, we provide a second mechanism for this purpose. This mechanism uses a linear programming techniques to compute the *fractional share* of a channel for each user, such that social welfare is maximized subject to each user receiving a share that is at least proportional to its max-min share in the conflict graph. We then adapt a decomposition technique due to Lavi and Swamy [17], to decompose the fractional shares to a set of feasible spectrum allocations. We show how to compute a probability distribution on this set, using a pair of primal and dual linear programs, such that picking a solution with the associated probability leads to the desired welfare-fairness trade-off between, in expectation.

In the rest of the paper, we present our system model and preliminaries in Sec. II. We design truthful auctions that strike a balance between social welfare and fairness for local fairness in Sec. III, and for max-min fairness in Sec. IV. We discuss related work in Sec. V, and conclude in Sec. VI.

II. PRELIMINARIES

We now present the system model in Sec. II-A, and provide some background in truthful auction design techniques in Sec. II-B.

A. System Model

We adopt the convention in auction theory and refer to secondary users as *agents*, and the primary license holder as the auctioneer. The set of agents is \mathcal{M} , and $|\mathcal{M}| = M$. Each agent $i \in \mathcal{M}$ is equipped with a single cognitive radio, capable of switching to different operating frequencies, or *channels*. Denote by $r(i)$ the transmission radius of the cognitive radio belonging to agent i , with $r(i)$ bounded for all agents i : $R_{min} \leq r(i) \leq R_{max}$. In practical settings, these ranges are somewhat similar, and the ratio $\frac{R_{max}}{R_{min}} = \Delta$ is a small constant. Two agents i and j *interfere* if both are assigned the same range of spectrum, and $d(i, j) \leq r(i) + r(j)$, where $d(i, j)$ is the spatial distance between their cognitive radios. We capture these interference constraints using a conflict graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \mathcal{M}$, and edge $(i, j) \in \mathcal{E}$ if radios of agents i and j potentially interfere. Since agents are also nodes in the conflict graphs, we will use terms agents and nodes interchangeably.

The auctioneer (primary spectrum user), owns large chunks of spectrum, whose utilization varies across time. As such, the auctioneer periodically pools unused spectrum together, which it divides into fixed sized blocks or *channels*, to be leased or assigned to secondary users. These leases are determined using *auctions*. We denote the set of channels as \mathcal{K} , with $|\mathcal{K}| =$

K. A *channel allocation* \mathbf{x} is an assignment of channels to agents, such that for each $i \in \mathcal{M}$, $x(i, k) = \{0, 1\}$ is a binary variable indicating whether i is assigned channel $k \in \mathcal{K}$. A channel allocation is *feasible* if no two agents that interfere are assigned the same channel. That is, for all channels k , if $x(i, k) = x(j, k) = 1$, then $(i, j) \notin \mathcal{E}$.

We assume that every agent $i \in \mathcal{M}$ is interested in obtaining at most one channel, and ascribes a certain value to being assigned one. Channels are indistinguishable, so that this value is the same for receiving any channel. The agent's valuation is assumed to be private information, known only to the agent itself. Let the valuation of agent i , v_i , be measured in monetary units. Spectrum auctions are held periodically. Each agent that wins a channel is given a lease for that channel, which expires at the beginning of the next round of the auction. At the start of each auction round, the auctioneer solicits *bids* from each agent. We denote agent i 's bid as b_i , and alternately use either \mathbf{b} or $(b_i \mathbf{b}_{-i})$ to indicate the vector of all bids, where \mathbf{b}_{-i} is the set of bids for all agents except i .

B. Strategyproof Auction Design

An auction can be viewed as a function that maps a given set of bids to (i) a channel allocation \mathbf{x} , and (ii) a payment vector \mathbf{p} , where $p(i)$ is the payment of each agent i . The utility of an agent is therefore a function of the set of submitted bids:

$$u_i(b_i, \mathbf{b}_{-i}) = \begin{cases} v_i - p(i) & \text{if } i \text{ receives a channel} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

In terms of strategic behaviour, we will adopt the usual convention in economics and assume that agents are *selfish* and *rational*, in the sense that their goal is to maximize (1).

We consider sealed-bid auctions, in which agents submit their bids at the beginning of each auction round. The auctioneer then determines the channel allocation \mathbf{x} , and the payment for all agents \mathbf{p} . Channel assignments expire at the beginning of the next auction round, and thus agents must submit a bid for each round in which it is interested in utilizing spectrum. Since agent valuations are private information, an agent may choose to submit a bid $b_i \neq v_i$, if this can lead to a higher utility. The auctioneer instead wishes to achieve an outcome that fits a pre-determined criteria such as maximizing social welfare, or achieving a fair allocation, or an outcome that is a tradeoff between both measures. In order to do, it is important for the auctioneer to elicit *truthful* bids from agents.

An auction is called *dominant-strategy truthful* if reporting the true valuation is the dominant strategy for an agent i , regardless of other agents' bids. More formally, an auction is dominant-strategy truthful or *strategyproof*, if for any agent i , for all $b_i \neq v_i$ and for any \mathbf{b}_{-i} , the following always holds in any auction outcome:

$$u_i(v_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i}) \quad (2)$$

Similarly, an auction is *truthful in expectation* if (2) holds in expectation. Besides the strategyproof property, we will

also require that the auction mechanisms we design fulfill the following as well:

$$u_i \geq 0, \quad \text{and } p_i \geq 0, \forall i \in \mathcal{M}$$

These two properties are known in the literature as *individual rationality* and *no positive transfers*, respectively [18]. The first ensures that agents do not suffer as a result of participating in the auction, while the second prevents the auctioneer from having to pay agents.

The best known mechanism for securing truthful bids from agents is the celebrated Vickrey-Clarke-Groves (VCG) mechanism [12], [14], [15]. Within the context of spectrum allocation, the VCG mechanism computes an optimal channel allocation \mathbf{x}^* , and charges each user the following

$$p(i) = \sum_{j \neq i} \sum_k v_j z^*(j, k) - \sum_{j \neq i} \sum_k v_j x^*(j, k) \quad (3)$$

where the allocation \mathbf{z}^* is computed by setting $v_i = 0$. It is easy to show that the payment scheme induces truthful behaviour, and we refer the reader to the text of Nisan *et al.* [18] for details.

Unfortunately, the VCG mechanism fails to be strategyproof when one does not have access to the optimal solution [16]. The problem of channel assignment is equivalent to graph colouring, which is NP-hard [19]. Hence, the auctioneer needs to resort to approximation algorithms in general. In Sec. III-B, we will design a truthful auction that employs an approximation algorithm to compute the channel allocation. Towards this direction, we will rely on Myerson's [20] characterization of truthful mechanisms.

Lemma 1. [Myerson, 1981] Let $P_i(b_i)$ be the probability of bidder i with bid b_i winning an auction. A mechanism is strategyproof if and only if the followings hold for a fixed \mathbf{b}_{-i} :

- $P_i(b_i)$ is monotonically non-decreasing in b_i
- Bidder i bidding b_i is charged $b_i P_i(b_i) - \int_0^{b_i} P_i(b) db$

There are two equivalent ways to interpret Lemma 1: (i) there exists a *minimum bid* b'_i such that i will win only if it bids at least b'_i , or (ii) the payment charged to i for a fixed \mathbf{b}_{-i} is independent of b_i . We will use the first point of view when designing a payment scheme for our auction in Sec. III-B.

III. A FRAMEWORK FOR TRUTHFUL AUCTIONS FOR BALANCING SOCIAL WELFARE WITH FAIRNESS

In this section, we will present our truthful auction framework for computing a channel allocation that provides a tradeoff between approximately optimal social welfare and fairness.

A. Social Welfare Maximization

A common goal of designing auctions is maximizing the *social welfare* [18], which is the total utility of all agents in the system, including the auctioneer. Since the auctioneer's utility is the sum of all payments received, one can state the social welfare simply as the total utility of all agents, $\sum_{i \in \mathcal{M}} \sum_{k \in \mathcal{K}} v_i x(i, k)$. We can state the optimal social welfare s as a function of the set of agent valuations, \mathbf{v} , the set

of available channels \mathcal{K} and the conflict graph \mathcal{G} . One can compute s using the following integer program (IP):

$$\begin{aligned} & \text{Maximize } s(\mathbf{v}, \mathcal{K}, \mathcal{G}) = \sum_{i \in \mathcal{M}} \sum_{k \in \mathcal{K}} v_i x(i, k) & (4) \\ & \text{Subject To:} \\ & x(i, k) + x(j, k) \leq 1 \quad \forall (i, j) \in \mathcal{E}, \forall k \in \mathcal{K} \\ & \sum_{k \in \mathcal{K}} x(i, k) \leq 1 \quad \forall i \in \mathcal{M} \\ & x(i, k) \in \{0, 1\} \quad \forall i \in \mathcal{M}, \forall k \in \mathcal{K} \end{aligned}$$

The problem of maximizing social welfare in this setting is equivalent to the well studied graph colouring problem, which is NP-Hard [19]. In practice, solving (4) optimally is infeasible, especially in the setting of a real-time secondary spectrum auction. We next design an auction that is truthful, but computes only an approximately optimal solution to the social welfare maximization problem.

B. A Truthful Auction for Approximately Max Social Welfare

Algorithm 1: A truthful auction for approximately maximizing social welfare with fairness constraints.

Input: Conflict graph \mathcal{G} , bid vector \mathbf{b} , set of random variables $\{X_i\}_{i \in \mathcal{M}}$

Output: Channel allocation \mathbf{x} , payment vector \mathbf{p}

```

1 foreach  $i \in \text{set} \mathcal{M}$  do
2    $x(i, k) := 0 \quad \forall i \in \mathcal{K}$ ;
3    $p(i) := 0$ ;
4    $\text{paid}(i) := \text{False}$ ;
5    $\text{saturated}(i) := \text{False}$ ;
6    $\phi(i) := \frac{X_i b_i}{N(i)+1} \quad \forall i \in \mathcal{N}$ ;
7   Let  $\mathcal{H} := \{\phi(i)\}_{i \in \mathcal{M}}$ ;
8   while  $\mathcal{H} \neq \emptyset$  do
9      $i := \arg \max_{\phi(i)} \{\mathcal{H}\}$ ;
10     $\mathcal{H} := \mathcal{H} \setminus \{i\}$ ;
11    if  $\text{saturated}(i) = \text{True}$  then
12       $\mathcal{C} := \{j | j \in N(i) \text{ and } \text{paid}(j) = \text{False} \text{ and}$ 
13         $\sum_k x(j, k) = 1\}$ ;
14       $\phi(l) := \min_{\phi(j)} \mathcal{C}$ ;
15       $p(l) := \phi(i) \left( \frac{N(l)+1}{X_l} \right)$ ;
16       $\text{paid}(l) := \text{True}$ ;
17    else if  $\sum_{k \in \mathcal{K}} (x(i, k) + \sum_{j \in N(i)} x(j, k)) = K$  then
18       $\text{saturated}(i) := \text{True}$ ;
19    else
20      foreach  $k \in 1 \dots K$  do
21        if  $\sum_{j \in N(i)} x(j, k) = 0$  then
22           $x(i, k) := 1$ ;
23          Break;
24  Return  $(\mathbf{x}, \mathbf{p})$ ;

```

Our auction mechanism is shown in Algorithm 1. The algorithm takes as input the conflict graph \mathcal{G} , the vector of bids \mathbf{b} solicited from agents interested in obtaining spectrum, as well as a set of random variables $\{X_i\}_{i \in \mathcal{M}}$. The random

variable X_i is defined for each agent i , and we will show in Sec. III-C how to set their values, to help achieve various objectives such as ensuring a particular form of fairness, or preventing agents from starving. The only requirement for X_i is that it remain *independent of* i 's bid, b_i . The auction begins by computing a *virtual bid* for each agent i , $\phi(i) = \frac{b_i X_i}{N(i)+1}$, which is proportional to the agent's bid and the value of X_i , but inversely proportional to its degree in the conflict graph. This term in the denominator helps us to rank high-degree agents lower when allocating channels. This idea is similar to the solution of Sakai *et al.* [21] for approximating a maximum weighted independent set. The auction then greedily assigns the available channels to agents, as long as it is feasible, in descending order of the virtual bid $\phi(i)$.

The allocation mechanism is simple, but we also need to ensure that our mechanism is able to compute payments for each winning agent that can guarantee truthful bidding. We will rely on Lemma 1 towards this direction. One way to interpret Myerson's lemma is to notice that that if the allocation rule posits a threshold bid b'_i , for which i is guaranteed to win as long as i bids $b_i \geq b'_i$, then charging i the price b'_i ensures that the mechanism is strategyproof. We refer the reader to the work by Archer and Tardos [22] for a concise treatment of Myerson's principle of truthfulness in one-dimensional settings. For now, let us focus on how one would be able to compute such a threshold bid given the allocation protocol used in Algorithm 1. Notice that since we greedily assign channels based on the virtual bid $\phi(i)$, the threshold bid of some winning agent i can be found by setting $\phi(i) = 0$ and running Algorithm 1 again with the rest of the input fixed as before. Then, the first agent j that gets assigned a channel such that it is no longer feasible to assign a channel to i is i 's *threshold agent*. Clearly, if $\phi(i) < \phi(j)$, then i would not be assigned a channel. We can then charge agent i the minimum bid required to ensure $\phi(i) \geq \phi(j)$. We can see that computing payments in this way will yield a strategyproof mechanism; we will make our argument more formally later.

The previous reasoning shows a simple recipe for computing the threshold bids for all agents, but it requires us to run the allocation protocol once to compute the channel allocation, and another $O(M)$ times to compute the payments for the winning agents. Such a solution is clearly undesirable in the context of real-time spectrum auctions that need to be light weight. Instead, we design a method for computing payments in a single run of the allocation algorithm, based on the following simple observation:

Lemma 2. Let $\phi(j)$ be the threshold virtual bid for a winning agent i in the outcome computed by Algorithm 1. We can claim that j is not assigned a channel in this outcome.

We say i is *saturated* when all K channels have been allocated to the set $\{i\} \cup \mathcal{N}(i)$. Lemma 2 shows that whenever we are unable to assign a channel to some agent i , then it must imply that i is a threshold agent for another agent l that was previously allocated a channel. Further, agent l 's payment must not have been previously computed, and l

must (i), be a neighbour of i , and (ii), be saturated. Therefore, in Algorithm 1, whenever we find a saturated agent i , we compute the agent l with the minimum virtual bid of all agents that fit requirements (i) and (ii), and use $\phi(i)$ to compute agent l 's payment in the following way

$$p(l) = \frac{N(l) + 1}{X_l} \phi(i) \quad (5)$$

We next show that our auction is individually rational, and truthful. We prepare by proving two lemmas first.

Lemma 3. The probability of i being allocated a channel is monotonically non-decreasing in its bid b_i . If an agent i is assigned a channel when bidding b_i , then it is also assigned a channel when bidding any $b'_i \geq b_i$.

This lemma is true since bidding higher can only increase an agent's expected virtual bid $\phi(i)$, and therefore increase its rank when being considered for allocation.

Lemma 4. The payment scheme of (5) is individually rational.

Proof: Since l wins when bidding b_l , it must be the case that $\phi(l) \geq \phi(i)$, where $\phi(i)$ is the threshold bid of agent l . This implies $b_l \geq \frac{b_i X_i}{N(i)+1} \frac{N(l)+1}{X_l} = p(l)$ ■

Theorem 1. The auction of Algorithm 1 is strategyproof.

Proof: Let v_i and b_i be agent i 's bid when being truthful and not truthful respectively, and let $x(v_i), x(b_i) \in \{0, 1\}$ be the allocation outcome, as a result of each bid. Fix all other agents' bids to be \mathbf{b}_{-i} . We examine the outcome of the auction on a case-by-case basis. First, assume that $b_i < v_i$. If $x(b_i) > x(v_i)$, then we get a contradiction due to Lemma 3. If $x(b_i) < x(v_i)$ or if $x(b_i) = x(v_i) = 0$, then clearly, i has no incentive to lie. If $x(b_i) = x(v_i) = 1$, then observe that the threshold bid is the same in both cases. Further, note that the random variables X_i and X_j are independent of i 's bid. Since i 's payment is given as $\frac{b_k X_k}{X_i} \frac{|\mathcal{N}(i)|+1}{|\mathcal{N}(k)|+1}$, i 's payment stays the same in both cases. Thus i once again has no incentive to lie. Now, assume that $v_i < b_i$ instead. If $x(b_i) < x(v_i)$, we get a contradiction due to Lemma 3. If $x(v_i) = x(b_i) = 0$, then i has no incentive to lie. The same is true for the case when $x(b_i) = x(v_i) = 1$, since the payment for agent i remains the same. Assume that $x(v_i) < x(b_i)$, then it means i is not allocated a channel when bidding truthfully. This suggests some neighbour k of i is allocated a channel, such that $\phi(k) \geq \phi(i)$, where $\phi(i)$ is i 's virtual bid when bidding truthfully. But this implies $v_i \leq \frac{b_k X_k}{N(k)+1} \frac{N(i)+1}{X_i}$. Since i 's threshold bid is at least $\phi(k)$, i 's payment when bidding b_i is also at least $\frac{b_k X_k}{N(k)+1} \frac{N(i)+1}{X_i} \geq v_i$, which means that i 's utility is at most 0. ■

C. Ensuring Fairness

In Sec. III-B, we designed a truthful auction framework through the use of virtual bids. The virtual bids were proportional to agents' true bids as well as the random variable X_i . We proved that such an auction can be made truthful and individually rational, as long as X_i was independent of the bid

vector. We now show how we can use these random variables to enforce 2-tier fairness, and prevent cognitive radios from starving. A key idea here is that the set of random variables $\{X_i\}_{i \in \mathcal{M}}$ provide the auctioneer with a way to *perturb the ranking* of the virtual bids, without breaking the strategyproof property of the auction.

Algorithm 2: Computing random variables $\{X_i\}_{i \in \mathcal{M}}$

Input: Set of agents \mathcal{M} , trade-off parameter ω , fairness function g

Output: Set of random variables $\{X_i\}_{i \in \mathcal{M}}$

```

1  $X_i := 0 \quad \forall i$ ;
2 Draw  $r$  uniformly at random from  $[0, 1]$ ;
3 if  $r < 1 - \omega$  then
4    $X_i := 1 \quad \forall i$ ;
5 else
6   Let  $\pi$  be random permutation of agents in  $\mathcal{M}$ ;
7   for  $j \in 1 \dots M$  do
8     if  $X_{\pi(j)} \neq 0$  then
9        $X_{\pi(j)} := g(\pi(j))$ ;
10      foreach  $k \in \mathcal{N}(\pi(j))$  do
11         $X_k := 1$ ;
12 return  $\{X_i\}_{i \in \mathcal{M}}$ ;

```

Algorithm 2 shows our algorithm for computing the random variables $\{X_i\}_{i \in \mathcal{M}}$. It takes as input the set of agents \mathcal{M} , a parameter ω that controls the trade-off between social welfare and fairness, and a function g for implementing different types of fair allocations. We will describe examples of g later. First, recall that the definition of two-tier fairness which attempts to guarantee a minimum service to all agents, before maximizing the performance of the system. In Algorithm 2, we use the parameter $0 \leq \omega \leq 1$ to control the trade-off between social welfare and fairness. If $\omega = 0$, then with probability 1, we set $X_i = 1$ for all agents i . This then reduces the mechanism of Algorithm 1 to an auction that maximizes social welfare, with no fairness constraint. For $\omega > 0$, the auction guarantees a return of $(1 - \omega)$ -fraction of the available social welfare, *in expectation*.

Let us next see how different functions g can achieve different notions of fairness for $\omega > 0$. The simplest function is one that ensures each agent receives a minimum share, regardless of its bid. Then, let V_{MAX} be the max bid possible by any agent. The following function achieves our objective:

$$g(i) = V_{MAX} \quad (6)$$

In Algorithm 2, each agent i has probability $\frac{1}{M}$ of appearing before its neighbours in the permutation π . In this event, we set $X_i = V_{MAX}$ and for each neighbour $j \in \mathcal{N}(i)$, we set $X_j = 1$, thereby ensuring any agent i is assigned a channel with probability at least $\frac{\omega}{M}$. Choosing a small ω , the primary user can ensure that no agent starves regardless of how small its bid is. Nonetheless, it is still in the best interest of every agent to bid truthfully, since bidding higher can only improve its chances of being assigned a channel.

Of course, the primary user may desire a more realistic notion of fairness. Going back to the example of Fig. 1, the primary user may only be interested in ensuring that agents who suffer from high interference be occasionally allocated a channel, as long as their valuation is at least as high as their neighbours. The following function can be used to achieve this objective:

$$g(i) = \frac{N(i) + 1}{2} \quad (7)$$

Observe that since every neighbour of i has degree at least 1, setting $X_i = \frac{N(i)+1}{2}$ guarantees that if an agent has valuation no lower than its neighbours', then it will receive a channel with probability at least $\frac{\omega}{M}$.

IV. MAX-MIN FAIR AUCTION

In this section, we will design a randomized auction that achieves a desired trade-off between maximizing social welfare, as well as max-min fairness, as a *global* fairness criterion. Max-min fairness is a common measure for achieving fairness when allocating resources in a network. We will use randomization as well as linear programming as our main tools. In order to simplify the notation and exposition in this section, we will assume $K = 1$, but note that it is straightforward to extend our technique for any K .

We start by examining how to compute max-min shares in a conflict graph. Let $\mathcal{C}(\mathcal{G})$ be the set of all cliques in \mathcal{G} . We can use the standard water-filling procedure shown in Algorithm 3 to compute the max-min share of each agent, $m(i)$, where $0 \leq m_i \leq 1, \forall i$. This algorithm is expensive, since it requires the computation of all cliques in the conflict graph. However, this can be justified in the sense that it is only executed once, prior to the auction taking place.

Algorithm 3: Water-filling type algorithm to compute base max-min share for each agent i in the interference graph \mathcal{G}

Input: Conflict graph \mathcal{G}

Output: Max-min share of agents m

- 1 Let $\mathcal{C}(\mathcal{G})$ be set of all cliques in \mathcal{G} ;
 - 2 For each agent i , let $c(i) \in \mathcal{C}$ be the largest clique containing i ;
 - 3 Initialize $\mathcal{I} := \mathcal{M}$;
 - 4 Initialize $m(i) := 0, \forall i \in \mathcal{M}$;
 - 5 **while** $\mathcal{I} \neq \emptyset$ **do**
 - 6 Increase $m(i)$ uniformly for all $i \in \mathcal{I}$ until for some $i, \sum_{j \in c(i)} m(j) = 1$;
 - 7 $\mathcal{I} := \mathcal{I} \setminus c(i)$;
 - 8 **return** y ;
-

Our goal is to compute a randomized channel allocation that achieves a trade-off between max-min fairness and social welfare maximization. Let $0 \leq \theta \leq 1$ be a tunable parameter that allows the auctioneer to achieve this tradeoff, stated in a more formal way below:

Definition 1. A randomized allocation rule is said to be ω -max-min fair, if it maximizes social welfare, subject to the constraint that for each agent i with max-min share $m(i)$, agent i 's probability of being assigned a channel is at least $\omega m(i)$.

Our definition of ω -max-min fairness is intended to ensure that agents do not suffer from starvation. All agents can be assigned a minimum probability of being allocated a channel. Setting $\omega = 1$ implies that the auctioneer neglects the goal of social welfare maximization completely. Similarly, setting $\omega = 0$ means that the auctioneer places no weight on max-min fairness, and is interested only in maximizing social welfare. Next, we show how to compute such an allocation, and then focus on designing a mechanism that can achieve such an allocation in a truthful fashion.

Naturally, *no feasible integral* channel assignment can meet the requirement of each agent having a minimum fractional channel allocation. Instead, we will first compute the desired fractional allocation. Later, we will show how to use a decomposition technique to build a set of feasible integral solutions, each associated with a probability, such that picking a feasible assignment with this probability leads to our desired goal of ω -max-min fairness in expectation. Consider the following linear program for computing a fractional allocation that maximizes social welfare, subject to each agent being assigned a minimum share:

$$\text{Maximize } f(\theta, \mathbf{v}) = \sum_i v_i x(i) \quad (8)$$

Subject To:

$$\begin{aligned} x(i) &\geq \theta m(i) && \forall i \in \mathcal{M} \\ x(i) + x(j) &\leq 1 && \forall (i, j) \in \mathcal{E} \\ x(i) + \sum_{j \in \mathcal{N}(i)} x(j) &\leq I(\mathcal{G}) && \forall i \in \mathcal{M} \\ x(i) &\geq 0 && \forall i \in \mathcal{M} \end{aligned}$$

The objective function of LP (8), $f(\theta, \mathbf{v})$ is a function of both the valuations of agents as well as the parameter $0 \leq \theta \leq 1$. Observe that for $\theta = 0$, maximizing $f(0, \mathbf{v})$ reduces to maximizing the fractional social welfare. Furthermore, since LP (8) is a maximization problem, $f(0, \mathbf{v}) \geq f(\theta, \mathbf{v})$ for any $\theta > 0$. The first constraint in LP (8) requires that each agent be assigned a minimum share proportional to its max-min share, $\theta m(i)$. The second constraint is sufficient to ensure a feasible solution. However, note that $f(0, \mathbf{v})$ without the third set of constraints is equivalent to the LP relaxation of the IP $s(\mathbf{v}, 1, \mathcal{G})$. Therefore, leaving out the third set of constraints defines a polytope such that the optimal solution for $f(0, \mathbf{v})$ has an integrality gap of $O(M)$ with respect to the optimal feasible integral solution of $s(\mathbf{v}, 1, \mathcal{G})$. To see this, consider the case when $\theta = 0$ for the complete graph with M nodes. It is known that the polytope of the LP relaxation of (4) always has a vertex solution that is half-integral [23]. In the complete graph, this leads to an optimal solution of $\frac{M}{2}$, obtained by setting $x(i) = \frac{1}{2}$ for all i . Since the optimal solution is 1, the integrality gap is $O(M)$. However, the performance of the auction we will design depends crucially on this integrality gap. Hence, we circumvent this problem by adding the third set of constraints to LP (8). For this set of constraints, $I(\mathcal{G})$ is a function of the conflict graph \mathcal{G} , that measures the size

of the maximum independent set consisting of the set of neighbours $\mathcal{N}(i)$ for any node i . Clearly, adding this constraint ensures that the polytope of LP (8) will continue to enclose all feasible integral channel assignment. The next two lemmas prove that (i) $I(\mathcal{G})$ is a constant in our network model, and (ii), the integrality gap is a constant that depends on the graph parameter Δ :

Lemma 5. For any node i , the size of the maximum independent set consisting of nodes in $\mathcal{N}(i)$, $I(\mathcal{G})$, is at most $\lfloor \frac{2\pi}{\arcsin(\frac{1}{2\Delta+1})} \rfloor - 1$.

Proof: We employ a geometric argument. For a given i , we need to find the max number of nodes that interfere with i but do not interfere with one another. The worst case occurs by letting $r(i) = R_{max}$, and placing n non mutually interfering nodes j with $r(j) = R_{min}$ on the circumference of a disk centered at i with radius $R_{max} + R_{min}$. Now, consider placing these n nodes at equidistance from each other along the circumference of this disk. We have

$$n = \frac{2\pi}{\arcsin(\frac{R_{min}}{2R_{max}+R_{min}})} = \frac{2\pi}{\arcsin(\frac{1}{2\Delta+1})}$$

The maximum number of nodes that can be place along the circumference without mutually interfering is $\lfloor n \rfloor - 1$, which yields the lemma. ■

Lemma 6. The integrality gap between $f(0, \mathbf{v})$ and $s(\mathbf{v}, 1, \mathcal{G})$ is $I(\mathcal{G}) = \lfloor \frac{2\pi}{\arcsin(\frac{1}{2\Delta+1})} \rfloor - 1$.

Proof: From Lemma 5, the size of the maximum independent set for any set of nodes $\{\{i\} \cup \mathcal{N}(i)\}$ is the constant $I(\mathcal{G})$, which is also the upper-bound of the constraint $x(i) + \sum_{j \in \mathcal{N}(i)} x(j)$. In the worst case, the integral solution picks only one node from this set (it must pick at least one). Since this is true for any node, and $x(i) \leq 1$ for all i , the lemma follows. ■

Next, we will employ a technique first shown by Lavi and Swamy [17] to decompose the fractional solution of LP (8) into a set of feasible solutions, \mathcal{L} . We denote by $l \in \mathcal{L}$ each feasible solution, such that for any $i, j \in l$, $(i, j) \notin \mathcal{E}$. For each solution $l \in \mathcal{L}$, we will find an associated probability $P(l)$, such that $\sum_{l \in \mathcal{L}} P(l) = 1$. Ideally, given an optimal solution x^* of LP (8), we would like to compute a probability distribution over feasible solutions such that $\sum_{l: i \in l} P(l) = x^*(i)$. This implies that picking a solution $l \in \mathcal{L}$ at random with probability $P(l)$ yields an allocation that is θ -max-min fair in expectation. Unfortunately, it turns out that this is not possible, since it would require us to compute all possible feasible channel assignments [17], which is NP-Hard. Instead, we will compute a set of probabilities such that for any i , $\sum_{l: i \in l} P(l) = \frac{x^*(i)}{C}$, where C is a small constant to be determined later. The decomposition technique relies on the ellipsoid algorithm [24] for solving linear programs. Denote by y^l a feasible channel allocation corresponding to the solution l . That is, $y^l(i) = 1$ means that in solution l , i is assigned a channel. Let x^* be the optimal solution to LP (8). Then the following linear program computes probability assignments $P(l)$ for each solution $l \in \mathcal{L}$, such that the probability of

agent i being assigned a channel is exactly $\frac{x^*(i)}{C}$:

$$\text{Minimize} \quad \sum_{l \in \mathcal{L}} P(l) \quad (9)$$

Subject To:

$$\begin{aligned} \sum_{l \in \mathcal{L}} x^l(i)P(l) &= \frac{x^*(i)}{C} \quad \forall i \in \mathcal{M} \\ \sum_{l \in \mathcal{L}} P(l) &\geq 1 \\ P(l) &\geq 0 \quad \forall l \in \mathcal{L} \end{aligned}$$

The dual of LP (9) is:

$$\text{Maximize} \quad \lambda + \sum_{i \in \mathcal{M}} \frac{x^*(i)}{C} \gamma(i) \quad (10)$$

Subject To:

$$\begin{aligned} \sum_{i \in \mathcal{M}} x^l(i) \gamma(i) &\leq 1 - \lambda \quad \forall l \in \mathcal{L} \\ \gamma(i) &\geq 0 \quad \forall i \in \mathcal{M} \end{aligned}$$

The variable $\gamma(i)$ in the dual corresponds to the first constraint of LP (9), while λ corresponds to the second. The decomposition technique will require the use of an algorithm to solve the maximum weighted independent set problem. We choose the polynomial time approximation scheme (PTAS) of Erlebach *et al.* [25], which we denote as \mathcal{A} . Algorithm \mathcal{A} computes a $(1 - \epsilon)$ approximation to the social welfare maximization problem, $f(0, \mathbf{v})$ for any set of valuations \mathbf{v} , for any $\epsilon > 0$. We will see later that the constant C in LP (9) depends on the approximation factor of \mathcal{A} , but for now, define $C = (1 - \epsilon)\alpha(\mathcal{G})$. We are now ready to show the following lemma, using complementary slackness conditions:

Lemma 7. The optimal solution to LP (9) is 1.

Proof: From the second constraint of LP (10), the solution is at least one. We next prove that the optimal solution (λ^*, γ^*) is at most 1 by way of contradiction. Assume this is not true for some instance of the LP. Then in the optimal solution, the second constraint of LP (9) is not tight. By complementary slackness, this implies $\lambda^* > 0$. Furthermore, by LP duality, we must have $\lambda^* + \sum_{i \in \mathcal{M}} \frac{x^*(i)}{C} \gamma^*(i) > 1$. If we treat the dual variable $\gamma^*(i)$ as the valuation of agent i , then, using the $(1 - \epsilon)$ -approximation algorithm \mathcal{A} , we can compute a social welfare maximizing feasible solution x^l , such that

$$\sum_i \gamma^*(i) x^l(i) \geq \frac{1}{1 - \epsilon} s(\gamma^*, 1, \mathcal{G}) \quad (11)$$

Further, from Lemma 6, we have

$$f(\theta, \gamma^*) \leq f(0, \gamma^*) \leq \alpha(\mathcal{G}) s(\gamma^*, 1, \mathcal{G}) \quad (12)$$

Substituting $C = (1 - \epsilon)\alpha(\mathcal{G})$, and using equations (11) and (12), we get

$$f(\theta, \gamma^*) = \sum_i \gamma^*(i) x^*(i) \leq C \sum_i \gamma^*(i) x^l(i)$$

Combined with $\lambda^* > 0$, this implies that $\sum_i \gamma^*(i) x^l(i) + \lambda^* \geq \frac{1}{C} \sum_i \gamma^*(i) x^*(i) + \lambda^* > 1$, *i.e.*, we have violated the second inequality of LP (10) for the solution x^l , contradicting the optimality of the solution (λ^*, γ^*) . ■

We have shown that the set $\{P(l)\}_{l \in \mathcal{L}}$ constitutes a valid probability distribution. However, solving LP (9) is difficult, primarily because we need to consider the variable x^l for all possible feasible solutions $l \in \mathcal{L}$, which is exponential in size. Instead of doing this directly, we will consider the dual LP (10), which has $M + 1$ variables, but an exponential number of constraints. The dual LP can be solved using the ellipsoid algorithm while employing the $(1 - \epsilon)$ -PTAS \mathcal{A} as a separation oracle [24]. The key insight here is that one can use a similar argument to the proof of Lemma 7, to show that for any solution where $\lambda + \sum_{i \in \mathcal{M}} \frac{x^*(i)}{C} \gamma(i) < 1$, there must be an integral solution x^l for which the constraints of LP (10) is violated. Further, we can use \mathcal{A} to find this solution. Since the ellipsoid algorithm is guaranteed to take at most a polynomial number of steps, it can be used to return a set of solutions $\{x^l\}_{l \in \mathcal{L}}$ that is polynomial in size. We can then plug these solutions back into LP (9) to compute the set of probabilities $\{P(l)\}_{l \in \mathcal{L}}$.

Algorithm 4: A truthful ω -max-min fair in expectation auction

Input: Conflict graph \mathcal{G} , bid vector \mathbf{b} , max-min shares $\{m(i)\}_{i \in \mathcal{M}}$, fairness parameter ω

Output: Channel assignment x^l , payment p

- 1 $\theta := C\omega$;
 - 2 Compute $f(\theta, \mathbf{b})$ using LP (8), let x^* be solution ;
 - 3 Use ellipsoid algorithm and \mathcal{A} on LP (10) with x^* to compute the polynomially sized set of solutions $\{x^l\}_{l \in \mathcal{L}}$;
 - 4 Solve LP (9) to compute probabilities $\{P(l)\}_{l \in \mathcal{L}}$;
 - 5 Pick some solution l' with probability $P(l')$;
 - 6 Initialize $p(i) := 0 \quad \forall i \in \mathcal{M}$;
 - 7 **foreach** i such that $x^l(i) = 1$ **do**
 - 8 Compute $f(\theta, (b_i = 0, \mathbf{b}_{-i}))$ with LP (8), let z be solution ;
 - 9 $p(i) := \frac{1}{x^*(i)} (\sum_{j \neq i} b_j z(j) - \sum_{j \neq i} b_j x^*(j))$;
 - 10 **return** (x^l, p) ;
-

The auction then works as shown in Algorithm 4. We first fix the parameter $\epsilon > 0$ for the PTAS \mathcal{A} , which gives us a trade-off between the running time of \mathcal{A} and the optimality of the solution obtained. We then set $\theta = C\omega$, where $C = (1 - \epsilon)\alpha(\mathcal{G})$. The fractional allocation $f(\theta, \mathbf{b})$ is next computed using LP (8). We apply the decomposition technique previously described to build a set of feasible solutions as well as an associated probability distribution. A channel allocation x^l is then picked probability $P(l)$. For each bidder i that is allocated a channel in x^l , we charge them the following:

$$p(i) = \frac{1}{x^*(i)} \left(\sum_{j \neq i} v_j z(j) - \sum_{j \neq i} v_j x^*(j) \right) \quad (13)$$

The solution z in (13) is obtained by computing the fractional allocation of LP (8) with $b_i = 0$. It can be shown that this payment scheme results in an auction that is truthful.

Theorem 2. Algorithm 4 is truthful in expectation.

Proof: Let $u_i(v_i)$ and $u_i(b_i)$ be the utility of agent i when bidding v_i and $b_i \neq v_i$ respectively. Similarly, let x and y be the solutions to LP (8) when i bids either v_i or b_i . Fix all other bids \mathbf{b}_{-i} . The expected utility of i when bidding truthfully is

$$\begin{aligned} E[u_i(v_i)] &= \frac{x(i)}{C} \left[v_i - \frac{1}{x(i)} \left(\sum_{j \neq i} b_j z(j) - \sum_{j \neq i} b_j x(j) \right) \right] \\ &= \frac{1}{C} \left[v_i x(i) + \sum_{j \neq i} b_j x(j) - \sum_{j \neq i} b_j z(j) \right] \end{aligned}$$

Since the polytope remains unchanged when the bid vector is (b_i, \mathbf{b}_{-i}) instead of (v_i, \mathbf{b}_{-i}) , the solution y is a feasible solution for the latter. This, together with the optimality of x for the bid vector (b_i, \mathbf{b}_{-i}) yields

$$E[u_i(v_i)] \geq \frac{1}{C} (v_i y(i) + \sum_{j \neq i} b_j y(j) - \sum_{j \neq i} b_j z(j)) = E[u_i(b_i)]$$

thus proving the theorem. \blacksquare

Algorithm 4 picks each agent i with probability $\frac{x^*(i)}{C}$, and from the first constraint of LP (8), it must be that $x^*(i) \geq \theta m(i) = C\omega m(i)$. Hence, each agent has a minimum probability of $\omega m(i)$ of being allocated a channel, which leads to the following:

Theorem 3. The mechanism shown in Algorithm 4 computes a ω -max-min fair allocation in expectation.

V. RELATED WORK

A plethora of recent studies show that the currently allocated spectrum sees usage that varies drastically, both geographically and temporally [2], [3]. For much of the time, licensed spectrum is idle [26]. With increased demand for spectrum due to the rapid growth of wireless applications, *dynamic spectrum access* has been proposed as a solution for this problem. Dynamic spectrum access schemes primarily rely on *cognitive radios* [26], [27], which are flexible wireless devices capable of switching its operating frequency. In the *uncoordinated* approach to dynamic spectrum sharing, cognitive radios are required to perform complex spectrum sensing operations in order to ensure that its transmissions do not interfere with the transmissions of the primary user [26], [28]. The *coordinated* approach on the other hand gives control of spectrum usage to the primary user directly. In such schemes, the primary user pools unused spectrum, which is then periodically leased to secondary users for short time durations using *spectrum auctions* [6], [8].

Auctions have long been used as a mechanism for distributing scarce resources or goods amongst competing users. An excellent treatment of auction theory can be found in the monograph of Krishna [29]. A celebrated result in auction theory is the VCG mechanism, which is due to the seminal series of work by Vickrey [12], Clarke [14] and Groves [15]. The VCG mechanism is the best known truthful mechanism, but crucially, loses this property when applied to suboptimal algorithms [16]. This renders VCG unsuitable for use in real-time secondary spectrum auctions, since computing the optimal interference-free channel allocation is NP-Hard.

Auctions for allocating spectrum for dynamic spectrum access has received considerable attention recently. Subramanian *et al.* [8] designed a greedy graph-colouring based algorithm for allocating spectrum while maximizing revenue in cellular networks. Revenue-maximization is also the objective of Gandhi *et al.* [7], who employed linear programming to model interference constraints. However, both studies ignore the possibility for strategic behaviour. In contrast, Jia *et al.* [9] design an auction for maximizing revenue, while also guaranteeing truthful bidding is a dominant strategy for bidders. Zhou *et al.* [4] also consider bidding behaviour of secondary users, and showed that due to interference, greedy channel assignment coupled with a VCG payment scheme yields a mechanism that is not strategyproof. Instead, they propose a greedy channel allocation scheme coupled with a tailored payment scheme to ensure truthful bidding is a dominant strategy of secondary users bidding for spectrum. Wu *et al.* [5] propose a semi-definite programming based solution instead for allocating channels, and design a mechanism that is not only truthful, but scalable and resistant to collusion. The above work all neglect the repeated nature of secondary spectrum auctions, and do not take fair allocation into consideration.

VI. CONCLUSION

Secondary spectrum auctions are a promising approach for spectrum sharing. In the interest of fairness, the outcome of these repeating auctions should strive to increase the diversity of users allocated spectrum, instead of greedily maximizing social welfare. This mitigates starvation among users, encourages users to continue taking part in the auction and reduces the incentive for vindictive bidding. In this paper, we provide two mechanisms that are provably strategyproof to achieve this goal, for local and global fairness respectively. Our first mechanism allows the primary user to ensure a minimum level of service to secondary users in the system, while the second mechanism computes an allocation that maximizes social welfare subject to max-min constraints in expectation. In the future, we plan to extend our work to consider additional measures of fairness, while considering more general models of the secondary spectrum market.

REFERENCES

- [1] P. Crampton, in *Handbook of Telecommunications Economics*, M. Cave, S. Majumdar, I. Vogelsang. (Eds.). Elsevier Science B.V., Amsterdam, 2002, pp. 605–639.
- [2] M. McHenry, P. Tenhula, D. McCloskey, D. Roberson, and C. Hood, “Chicago spectrum occupancy measurements & analysis and a long-term studies proposal,” in *Proceedings of the ACM International Workshop on Technology and Policy for Accessing Spectrum*, 2006.
- [3] “Spectrum occupancy measurement reports. available online.” <http://www.sharespectrum.com/measurements/>.
- [4] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, “eBay in the Sky: strategy-proof wireless spectrum auctions,” in *Proceedings of ACM MOBICOM*, 2008.
- [5] Y. Wu, B. Wang, K. Liu, and T. Clancy, “A scalable collusion-resistant multi-winner cognitive spectrum auction game,” *IEEE Trans. on Communications*, vol. 57, no. 12, pp. 3805–3816, 2009.
- [6] “Spectrum bridge.” <http://www.spectrumbridge.com/>.

- [7] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri, “A general framework for wireless spectrum auctions,” in *Proceedings of the IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2007.
- [8] A. Subramanian, M. Al-Ayyoub, H. Gupta, S. Das, and M. Buddhikot, “Near-Optimal Dynamic Spectrum Allocation in Cellular Networks,” in *Proceedings of the IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2008.
- [9] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, “Revenue generation for truthful spectrum auction in dynamic spectrum access,” in *Proceedings of ACM MOBIHOC*, 2009.
- [10] C. Meek, D. Chickering, and D. Wilson, “Stochastic and contingent-payment auctions,” in *1st Workshop on Sponsored Search Auctions*, 2005.
- [11] Y. Zhou and R. Lukose, “Vindictive bidding in keyword auctions,” in *Proceedings of the International Conference on Electronic Commerce (ICEC)*, 2007.
- [12] W. Vickrey, “Counterspeculation, auctions, and competitive sealed tenders,” *Journal of finance*, pp. 8–37, 1961.
- [13] A. Goldberg and J. Hartline, “Envy-free auctions for digital goods,” in *Proceedings of the 4th ACM conference on Electronic commerce*. ACM, 2003, p. 35.
- [14] E. Clarke, “Multipart pricing of public goods,” *Public choice*, vol. 11, no. 1, pp. 17–33, 1971.
- [15] T. Groves, “Incentives in teams,” *Econometrica: Journal of the Econometric Society*, pp. 617–631, 1973.
- [16] N. Nisan and A. Ronen, “Computationally feasible VCG mechanisms,” *Journal of Artificial Intelligence Research*, vol. 29, pp. 19–47, 2007.
- [17] R. Lavi and C. Swamy, “Truthful and Near-Optimal Mechanism Design via Linear Programming,” in *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*, 2005.
- [18] N. Nisan, T. Roughgarden, E. Tardos, and V. V. (Eds.), *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [19] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-completeness*. W.H. Freeman and Company, San Francisco, 1979.
- [20] R. Myerson, “Optimal auction design,” *Mathematics of operations research*, pp. 58–73, 1981.
- [21] S. Sakai, M. Togasaki, and K. Yamazaki, “A note on greedy algorithms for the maximum weighted independent set problem,” *Discrete Applied Mathematics*, vol. 126, no. 2-3, pp. 313–322, 2003.
- [22] A. Archer and E. Tardos, “Truthful mechanisms for one-parameter agents,” in *Proceedings of IEEE Symposium on Foundations of Computer Science (FOCS)*, 2001.
- [23] A. Schrijver, *Combinatorial optimization: polyhedra and efficiency*. Springer Verlag, 2003.
- [24] C. Papadimitriou and K. Steiglitz, *Combinatorial optimization: algorithms and complexity*. Dover Publications, 1998.
- [25] T. Erlebach, K. Jansen, and E. Seidel, “Polynomial-time approximation schemes for geometric graphs,” in *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2001.
- [26] I. Akyildiz, W. Lee, M. Vuran, and S. Mohanty, “NeXt generation/dynamic spectrum access/cognitive radio wireless networks: a survey,” *Computer Networks*, vol. 50, no. 13, pp. 2127–2159, 2006.
- [27] J. Mitola, “Software radio architecture: a mathematical perspective,” *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 514–538, 1999.
- [28] Z. Ji and K. Liu, “Dynamic Spectrum Sharing: A Game Theoretical Overview,” *IEEE Communications Magazine*, vol. 45, no. 5, pp. 88–94, 2007.
- [29] V. Krishna, *Auction theory*. Academic Press, 2002.