Revenue-maximizing and Truthful Online Auctions for Dynamic Spectrum Access

Ajay Gopinathan[†], Niklas Carlsson[§], Zongpeng Li[†], Chuan Wu[‡]
[†] University of Calgary, Canada. {ajay.gopinathan, zongpeng}@ucalgary.ca
[§] Linköping University, Sweden. niklas.carlsson@liu.se
[‡] University of Hong Kong, China. cwu@cs.hku.hk

***Abstract*—Secondary spectrum auctions have been suggested as a strategically robust mechanism for distributing idle spectrum to competing secondary users. However, previous work on such auction design have assumed a static auction setting, thus failing to fully exploit the inherently time-varying nature of spectrum demand and utilization. In this paper, we address this issue from the perspective of the primary user who wishes to maximize the auction revenue. We present an online auction framework that dynamically accepts bids and allocates spectrum. We prove rigorously that our online auction framework is truthful in the multiple dimensions of bid values, as well as bid timing parameters. To protect against unbounded loss of revenue due to latter bids, we introduce controlled preemption into our mechanism. We prove that preemption, coupled with the technique of inflating bids artificially, leads to an online auction that guarantees a $\frac{1}{5}$-fraction of the optimal revenue as obtained by an offline adversary. Since the previous guarantee holds only for the optimal channel allocation, we further provide a greedy channel allocation scheme which provides scalability. We prove that the greedy scheme also obtains a constant competitive revenue guarantee, where the constant depends on the parameter of the conflict graph.**

I. Introduction

In recent years, government agencies have turned to the use of large-scale *auctions* when allocating newly available spectrum [1]. The winners of these auctions are known as *primary users*, and often hold exclusive spectrum usage rights for long time periods. The drawback of this static allocation model is that smaller entities end up starving for spectrum [1]–[3]. Significant variations in the primary users' spectrum utilization in both space and time [4] has motivated the creation of a *secondary spectrum market* [5].

In the dynamic spectrum access model, the primary user divides the available spectrum into fixed size chunks, or *channels*, to be leased to *secondary users* for a short period of time [2], [3], [6]. *Auctions* are a natural, revenue-generating mechanism to consider in this setting. Auctions adhere to well defined notions of fairness and economic feasibility, and have been shown to be an effective mechanism for coordinating secondary spectrum access [2], [3], [6]–[8]. More importantly, auctions are theoretically robust to strategic manipulation by bidders, and can be designed to guarantee truthful bidding behavior. Such auctions are said to be *strategyproof* [9], [10].

In contrast to traditional auctions, secondary spectrum auctions are unique in that they have distinct *spatial* and *temporal* characteristics. Spectrum offers the opportunity for *spatial reuse*. Multiple, secondary users can be sold the same range of spectrum as long as they do not interfere. Since spectrum demand by secondary users is likely to fluctuate in time, an *online auction* is more desirable than an *offline auction*. Offline auctions compute inflexible, fixed period allocations, thereby failing to fully exploit the opportunity afforded by the time-varying nature of spectrum demand. In contrast, an online auction allocates spectrum *on demand*, which improves the efficiency of spectrum usage [6], [11], while leading to potentially higher revenue for the primary user.

An online secondary spectrum auction presents a number of key design challenges. The auction needs to efficiently compute a feasible, interference-free channel allocation, maximize the revenue generated, and simultaneously remain strategyproof. In the online setting, achieving the strategyproof property requires new techniques. The direct application of known strategyproof mechanisms, like the Vickrey-Clarke-Groves (VCG) [9], is insufficient to ensure truthful bidding. The reason for this is that, in contrast to traditional, static auctions, the temporal aspect of online auctions presents the secondary user with an *additional* dimension for strategic manipulation. In particular, bidders may choose to not only manipulate their *valuations*, but also the *timing* of their bids if doing so improves their utility. Therefore, online auctions need to be strategyproof in *both* valuations as well as in the timing of the bids.

Another design challenge is presented by the need for online auctions to make allocation decisions in real-time, without access to the set of bids that could arrive in the future. If the decision to allocate a channel is irrevocable, then *no online auction* can guarantee a constant fraction of the optimal revenue, as compared to an *offline adversary* with knowledge of future bids. The solution to this lies in *preemption* [6], [12]. Preemption allows the auction to place an upper bound on the loss of revenue due to allocation decisions that are, in hindsight, suboptimal.

In this paper, we address the above challenges and design an online auction framework for secondary spectrum access that maximizes revenue while remaining strategyproof in both valuations and timing of the bids. Our auction framework runs in real-time, and caters to time-varying demands by allowing secondary users to request for spectrum at any time. Our framework uses preemption to judiciously select winning bidders, which yields a mechanism with a constant approximation revenue guarantee. Our framework is flexible, and can be used in conjunction with the algorithm that computes the optimal channel allocation scheme, or even one that only finds an approximately optimal allocation through greedy assignment. In both cases, the mechanism is tailored to ensure strategyproofness. When the optimal allocation algorithm is used, we show that our online framework recovers a $\frac{1}{5}$-fraction of the optimal revenue, and that this bound is tight. The VCG mechanism is not strategyproof when applied directly. Instead, we show that the VCG payment must be computed for every bidder in *each* timeslot for which the bid is valid. We show that choosing the *minimum* payment from this schedule is sufficient to ensure bidders have no incentive to delay their bids.

Since the channel assignment problem is equivalent to the NP-Hard problem of graph coloring [13], we further design an efficient greedy channel allocation scheme. We show that this algorithm too, is constant competitive with respect to the optimal revenue. Our algorithm also reduces unnecessary preemptions compared to the naive greedy algorithm. Since the allocation computed is suboptimal, using the VCG mechanism to compute payments does not guarantee truthful bidding [14]. Hence, we tailor a payment scheme based on the idea of threshold bids to recover the strategyproof property.

This paper is organized as follows. Section II discusses related work. We then introduce our system model in Section III. We present our auction framework in Section IV, and analyze the performance of this framework for the optimal and greedy channel allocation schemes in Sections V and VI respectively. Section VII summarizes our work.

II. Related Work

In economic theory, auctions are commonly used for allocating scarce resources amongst competing users. Auctions are especially attractive since they are provably robust to strategic bidding behavior [15], [16]. The best known strategyproof auction is the Vickrey-Clarke-Groves (VCG) mechanism [9]. In this paper, we employ the VCG payment scheme as a component to ensure our online auction framework is truthful when computing the optimal spectrum allocation.

There are a number of drawbacks to using the VCG auction scheme. First, VCG auctions in general have poor revenue generating properties [9]. Second, VCG is not truthful when used in conjunction with suboptimal solutions [14]. This presents a problem, since computing the optimal, interference-free allocation is NP-Hard [13]. For an online auction to be truly scalable, approximation algorithms that run in polynomial time must be used instead. In this paper, we solve both these problems by leveraging the seminal work of Myerson [17]. Myerson's work not only provides the foundation for revenue-generating auctions, but also gives a mathematical characterization for all truthful mechanisms. We use this characterization to build a greedy, approximately optimal channel allocation scheme that remains truthful without resorting to the VCG mechanism.

Several recent studies have looked at employing auctions for dynamic secondary spectrum access [2], [3], [6], [7], [11], [18]. In particular, the work of Zhou *et al.* [3] showed how to design truthful auctions for maximizing social welfare while exploiting the potential for spatial reuse of channels. In contrast, Jia *et al.* [2] and Gopinathan and Li [7] focus on maximizing revenue instead. All of the previously mentioned work focus on the offline setting. In contrast, we focus on the online setting in this paper for dynamic, real-time spectrum allocation. Within this setting, Xu *et al.* study spectrum admission in the case where bid arrivals are assumed to follow a Poisson distribution [11]. However, they do not consider the possibility of preemption and the reuse of spectrum in the time domain. In contrast, Deek *et al.* [6] propose an online auction which allows preemption, and is strategyproof in the time domain. Both these auctions focus on maximizing social welfare, which in general does not guarantee good revenue. In contrast, our online auction framework is explicitly focused on revenue generation in a truthful manner, and we provide theoretical bounds on the amount of revenue obtained as compared to the optimal, offline algorithm.

III. Preliminaries

A. System Model

Let $\mathcal{M} = \{1, 2, \dots, M\}$ denote the set of secondary users. As is customary in auction theory, we will occasionally refer to secondary users as *agents*. Each secondary user i is assumed to be equipped with a cognitive radio that is capable of operating on different spectrum frequencies, or *channels*. Let $r(i)$ be the transmission radius of i's radio, such that $R_{min} \leq r(i) \leq R_{\max}$, where R_{min} and R_{max} are the minimum and maximum possible transmission radius. In practice, secondary users are fairly homogeneous, so that $\frac{R_{max}}{R_{min}} = \Delta \geq 1$ is a small constant. Secondary users $i, j \in \mathcal{M}$ are said to *interfere*, if they simultaneously transmit on the same frequency and the spatial distance $d(i, j)$ between them is less than the sum of their transmission radius, *i.e.*, $d(i, j) \leq r(i) + r(j)$. The conflict graph $\mathcal{G} = (\mathcal{M}, \mathcal{E})$, is used to model these interference constraints. We will denote using $\mathcal{N}_i$ the set of neighbors of i in $\mathcal{G}$.

Time in the system is divided into slots of equal length. In practice, the optimal length of the timeslot is a function of how the primary user's utilization and the demand by secondary users vary in time. The primary user divides the available spectrum into a set of K channels, $\mathcal{K}$. Each channel $k \in \mathcal{K}$ is assumed to be homogeneous. Demand for these channels is assumed to be dynamic and time-dependent. Each secondary user requires the use of any one of these K channels for T continuous timeslots, and is willing to pay for it.

For user i, let v_i be the *valuation* of i, which measures in monetary units the maximum amount that i is willing to pay for the exclusive use of the *same* channel for T continuous timeslots. While it is possible to relax this assumption, we will focus on the technically more challenging setting when i's demand is binary – agents have a value only for receiving the same channel for T continuous timeslots, and being assigned different channels for this period, or holding the channel for less than T timeslots, is assumed to have no value. Valuations are considered *private* information, known only to i and unknown to the primary user.

We will use the binary variable $x_i^k(t)$ to indicate that i is assigned the channel k in timeslot t. The vector $\mathbf{x}_i(t)$ represents the channel assignment for i at time t. A channel allocation for the entire system, described by the vector $\mathbf{x}(t)$ is said to be feasible or interference-free, if no two agents interfere during timeslot t.

B. Auction Format

In our model, secondary users arrive and depart the system dynamically, depending on when they require spectrum, and when they are granted access by the primary user. Without loss of generality, we assume arrival and departure events happen at the boundary of timeslots. At the beginning of every timeslot, zero or more users may submit a *bid* requesting spectrum from the primary user. The bid of user i, θ_i, takes the form

$$\theta_i = (v_i, a_i, d_i). \tag{1}$$

Here, a_i is the time when user i realizes its need for spectrum, and d_i is the deadline by which this demand must be satisfied. Without loss of generality, let $a_i \leq d_i$, and $a_i + T - 1 \leq d_i$. We adopt the convention in auction theory and assume that for each $i \in \mathcal{M}$, v_i is drawn independently and at random from the cumulative probability distribution F_i, where

$$F_i(v') = Prob(v_i \leq v').$$

As previously stated, the true value of the bid θ_i is private known only to i. However, the prior distribution of bids, $\mathbf{F} = F_1 \times F_2 \times \ldots \times F_M$ is public information known to the primary user. For example, the distribution $\mathbf{F}$ could have been learned via market research, or inferred from the prior bidding behavior of secondary users [9], [15].

Mathematically, an auction can be viewed as a function, $\mathcal{A}$, that maps the set of submitted bids to a tuple consisting of (i) a channel allocation $\mathbf{x}$, and (ii) a payment vector $\mathbf{p}$. The goal of the primary user is to compute $\mathbf{x}$ and $\mathbf{p}$ such that the sum of the latter is maximized while ensuring the former is interference free. The nature of secondary users that arrive and depart the system at various points in time contributes greatly to the difficulty of this problem. It is not possible for the auctioneer to compute $\mathbf{x}$ and $\mathbf{p}$ in advance. Instead, the auctioneer must recompute the auction at *every timeslot*, incorporating new information about secondary users and their bids, while taking into consideration the channel assignment already in place. Doing so in an optimal fashion requires an online auction algorithm, which is what we will provide in this paper.

When measuring the performance of our auction, we will use the *offline adversary* [19] as our benchmark. Such an adversary is assumed to have perfect knowledge of the sequence of future bid arrivals, and is therefore able to schedule the allocation optimally. The revenue approximation guarantees therefore hold for a *worse case setting*. Despite this, we will show that our auctions are within a constant factor of the optimal, even for such *worse case setting*.

C. Strategyproof Auction Design

Let $p_i(\theta_i, \boldsymbol{\theta}_{-i})$ be the payment charged to an agent who is allocated a channel when bidding θ_i while all other agents bid $\boldsymbol{\theta}_{-i} = \{\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots \theta_M\}$. The *utility* of the agent, $u_i(\theta_i)$, can be stated as

$$u_i(\theta_i, \boldsymbol{\theta}_{-i}) = \begin{cases} v_i - p_i(\theta_i), & \text{if } i \text{ is assigned a channel} \\ & \text{for } T \text{ continuous timeslots} \\ 0 & \text{otherwise} \end{cases}. \tag{2}$$

Agents are *selfish* and *rational*, and strategically may choose to misreport or manipulate their bid if doing so allows them to improve their utility [9], [15]. In our setting, an agent can lie not only about their valuation v_i, but may also choose to misreport the value d_i, or delay submitting her bid (*i.e.* lie about a_i). Let $\theta'_i = (v'_i, a'_i, d'_i)$ be the bid submitted by i (not necessarily truthfully). Similar to previous work in the area of online auctions [6], [11], [12], [20], we assume *no early arrivals* and *no late departures*. That is, $a'_i \geq a_i$ and $d'_i \leq d_i$.

It is imperative that the auctioneer (the primary user in our setting), obtains bids truthfully so the goal of maximizing revenue can be achieved. An auction is said to be *dominant strategy truthful*, or strategyproof, if for any agent i, for all $\theta'_i \neq \theta_i$, and for *any* $\boldsymbol{\theta}_{-i}$, it is always the case that $u_i(\theta_i, \boldsymbol{\theta}_{-i}) \geq u_i(\theta'_i, \boldsymbol{\theta}_{-i})$. The following characterization of strategyproof mechanisms by Myerson will provide the foundation for our approach.

Lemma 1. [Myerson, 1981] Let $x_i(b_i)$ be the allocation function used for bidder i with bid b_i. A mechanism is strategyproof if and only if, for any fixed set of bids by all other agents not including i, the following conditions hold:

- $x_i(b_i)$ is monotonically non-decreasing in b_i
- Bidder i bidding b_i is charged a payment computed as

$$b_i x_i(b_i) - \int_0^{b_i} x_i(b) db. \tag{3}$$

Lemma 1 states that in order to ensure a mechanism is truthful, one must use an allocation rule that is monotonic. Intuitively, this means that bidding higher should never harm the bidder's chances of winning the auction. Furthermore, in the special case when the allocation function is deterministic, it can be easily verified that the second condition in Lemma 1 requires that a winning bidder is charged the *minimum bid* that guarantees being allocated in the auction. Based on these observations, Lemma 1 provides a general purpose recipe for designing truthful auctions:- Design an approximation algorithm that is *monotone* in the input values, and ensure that each bidder is charged the *minimum bid* that guarantees allocation. We will term every mechanism that fits this bill as a *monotone mechanism*. In our setting, a bid spans the multiple dimensions of both valuation as well as time. Keeping this in mind, we design the online auction framework to be monotone in *both* of these domains.

D. Maximizing Revenue

An auction that maximizes revenue or profit for the auctioneer is also known as an *optimal auction* in economic theory. The problem of designing optimal auctions was first studied by Myerson [17]. Myerson introduced the notion of a *virtual valuation*, $\phi_i(v_i)$ of bidder i, where

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}. \tag{4}$$

Here, $f_i = \frac{dF_i}{dv}$ is the probability density function of F_i. Myerson's seminal work proved the following theorem.

Theorem 1. [Myerson, 1981] Given the valuation distribution $\mathbf{F} = F_1 \times F_2 \times \ldots \times F_M$, and truthful bids $\{v_1, v_2, \ldots, v_M\}$ such that v_i is drawn independently and at random from F_i

for all $i \in \mathcal{M}$, the expected revenue of *any* mechanism, for allocation of goods $\mathbf{x} = \{x_1, x_2, \ldots, x_M\}$ is given by

$$R = \sum_{i \in \mathcal{M}} \phi_i(v_i) x_i. \tag{5}$$

That is, in order to maximize revenue, (5) must be maximized. This is a rather powerful result, since it means that within the large space of all possible revenue maximizing mechanisms, we can restrict our focus on *truthful* mechanisms that optimize (5). We have already seen that any truthful mechanism requires an allocation function that is monotone. This therefore imposes an additional condition on $\mathbf{F}$. In particular, we require F_i to be *regular* for all agents i, in that $\phi_i(v_i)$ is monotone non-decreasing in v_i. This is in fact a rather mild assumption, since most natural distributions of interest (*e.g.*, uniform, exponential, Gaussian, etc.) are regular. Furthermore, there are readily available techniques for dealing with non-regular distributions [17].

IV. A Strategyproof Revenue-Maximizing Online Auction Framework

In this section, we will design a general purpose strategyproof online auction framework with the goal of maximizing revenue. The framework solicits bids at the beginning of each timeslot, and executes a monotone channel allocation and payment mechanism $\mathcal{A}$. Sections V and VI will describe two possible choices for mechanism $\mathcal{A}$, one that computes the optimal channel allocation, and another based on a greedy heuristic. While the framework puts into effect the allocation computed by $\mathcal{A}$ at every timeslot, the payment computed for the allocation is not used immediately. Instead, it records these payments for each timeslot in a *payment schedule*. We will show that by charging bidders the *minimum* payment computed in the schedule as the final price, the framework is strategyproof in time. The framework allows for the *preemption* of assigned channels in order to guarantee a constant approximation to the optimal revenue. We therefore begin by introducing the idea of inflating bids over time, which allows us to control the amount of preemption that takes place.

A. Preemption and Bid Inflation

An allocation decision made during a timeslot could have a severe impact on the revenue generated, especially if these decisions are *irrevocable*. Motivated by this, we allow preemption in our online auction framework. Since preemption can potentially be unfair to secondary users, our auction does not charge payment to bidders who do not receive the same channel for T continuous timeslots.

It is also crucial that we do not end up preempting unnecessarily. Intuitively, an agent that has completed $T-1$ timeslots should be harder to preempt than one who has only been assigned for a single timeslot so far. This naturally leads to the idea of *artificially inflating bids*, in proportion to the number of timeslots the bidder has continuously held the channel. In particular, if at time t, bidder i has held the channel k for $\tau_i < T$ continuous timeslots, then our online auction sets the virtual valuation ϕ_i as

$$\phi_i(v_i, \tau_i) = (1+\delta)^{\frac{\tau_i}{T}} \left(v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right), \tag{6}$$

Algorithm 1: The online auction framework executed for every timeslot t

Input: Conflict graph $\mathcal{G}$, bid pool Θ, set of channels $\mathcal{K}$, monotone allocation and payment mechanism $\mathcal{A}$, previous channel assignment $\mathbf{x}(t-1)$, current timeslot t
Output: Channel assignment $\mathbf{x}(t)$, payment schedule $\mathbf{s}(t)$

foreach $\theta_i \in \Theta$ **do**
  if $x_i(t-1) > 0$ **then**
    $\tau_i := \tau_i + 1$
    if $\tau_i = T$ **then**
      Require i to release channel
      $\Theta := \Theta \setminus \{\theta_i\}$
  else
    $\tau_i := 0$
Let $\mathcal{M}'$ be the set of agents with bids in $\mathbf{\Theta}$
Let $\mathcal{G}'$ be the subgraph induced by $\mathcal{M}'$
foreach $\theta_i \in \Theta$ **do**
  $\phi_i := (1+\delta)^{\frac{\tau_i}{T}} \left(v_i - \frac{1-F_i(v_i)}{f_i(v_i)} \right)$
Run $\mathcal{A}$ on $\mathcal{G}'$ using the set of bids $\{\phi_i\}_{i \in \mathcal{M}'}$, the previous channel allocation $\mathbf{x}(t-1)$, and the length of channel leases so far $\{\tau_i\}_{i \in \mathcal{M}'}$
Let $\mathbf{x}(t)$ be the channel allocation, and $\mathbf{p}(t)$ be the corresponding payment returned by $\mathcal{A}$
foreach $i \in \mathcal{M}$ **do**
  if $i \in \mathcal{M}'$ **then**
    $s_i(t) := p_i(t)$
  else
    $s_i(t) := \infty$
Return $(\mathbf{x}(t), \mathbf{s}(t))$

where $\delta > 0$ is a parameter that can be adjusted by the auctioneer. Observe that bids are now inflated by an *exponential* function with respect to the ratio of the number of continuous timeslots a user has been allocated to the total demand T. This idea was first introduced by Hajiaghayi *et al.* [12]. Intuitively, this allows us to ensure that the cost of a continuous sequence of preemptions ends up being a telescoping sum. Later, we will be more precise with this intuition, and show that bid inflation using (6) with $\delta = 1$ is sufficient to guarantee constant fraction of the optimal revenue with respect to the offline adversary.

B. The Online Auction Framework

Algorithm 1 describe our online auction framework in detail. The algorithm is executed at the beginning of each timeslot t. At this time, zero or more bids may be received by the primary user, which are then added into a *bid pool*, $\mathbf{\Theta}$. An element $\Theta_i \in \mathbf{\Theta}$ of the bid pool is a tuple (θ_i, τ_i), where θ_i is the bid submitted by bidder i, and τ_i is the number of continuous timeslots for which i has held a channel at time t. Initially, $\tau_i = 0$. The element Θ_i is removed from the pool only when $\tau_i = T$; that is, the user has been allocated a channel for T continuous timeslots. Also, Θ_i will be removed if the bid has expired; *i.e.*, at time $t = d_i - T$. A secondary user who has previously been assigned a channel may be preempted before her lease expires. If a user is preempted at time t, then the bid is returned to the pool if it is still valid; *i.e.*, $t < d_i - T$, otherwise, the bid is rejected.

At the beginning of every timeslot, bids in the bid pool are used as input to the channel allocation and payment mechanism $\mathcal{A}$. This algorithm is used to compute the channel allocation,

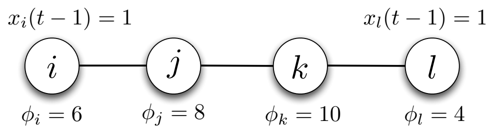

Fig. 1. The optimal assignment at time t must preempt agent l

as well as the *payment* $p_i(t)$, which is specific to timeslot t. This computed payment is not charged to agents immediately. Instead, we use it to update the *payment schedule* $\mathbf{s}$. Let $s_i(t)$ be the computed payment for user i at timeslot t, then the *final price* charged to a winning user is computed as

$$\min_{a_i \leq t \leq d_i - T + 1} s_i(t). \tag{7}$$

Thus, a winning user will always be charged the minimum price over the period during which the user's bid was valid.

Our auction framework leaves some flexibility for the primary user to choose the allocation and payment mechanism $\mathcal{A}$. We require only that $\mathcal{A}$ is monotone in the sense of Lemma 1. The performance and truthfulness of the auction framework is closely tied to the choice of $\mathcal{A}$. In the sequel, we analyze the performance of Algorithm 1 for when we use the optimal channel allocation mechanism, and for when we use a greedy channel allocation mechanism.

V. The Optimal Mechanism

A. The Optimal Channel Allocation with Preemption

Since our goal is to compute the optimal channel allocation, it is worth pausing for a moment to consider what exactly constitutes an optimal allocation given that channels can be preempted. Ideally, we would like to choose a channel assignment $\mathbf{x}(t)$ at time t that maximizes

$$\sum_{i \in \mathcal{M}'} \phi_i(v_i, \tau_i) x_i(t),$$

where $\mathcal{M}'$ is the set of agents with currently valid bids. From Theorem 1, we know that any truthful mechanism that does so also maximizes the expected payment, and hence revenue. However, assigning different channels to a user in consecutive timeslots counts as preemption in our model. Therefore, the optimal allocation *maximizes* the total value of the bidders assigned a channel, while *minimizing* the *cost* of preemption.

As an example, consider the network shown in Figure 1. In this network, there are only two channels, and bids during the current timeslot t are shown next to each user. The secondary users i and l were previously assigned channel 1, while users j and k submit new bids at time t. When computing the new allocation, it is preferable to preempt l instead of i to minimize the preemption cost. The optimal allocation is therefore

$$x_i(t) = 1 \qquad x_j(t) = 2 \qquad x_k(t) = 1 \qquad x_l(t) = 2.$$

The optimal allocation at time t can be computed with an integer linear program. The program is a function of the bids of the users in the system, suitably inflated according to (6) for users that already had a channel assigned at time $t-1$.

$$\text{Maximize} \quad \sum_{i \in \mathcal{M}} \phi_i(v_i, \tau_i^k) x_i^k(t) \tag{8}$$

Algorithm 2: The optimal allocation algorithm with a VCG payment scheme.

Input: Induced conflict graph $\mathcal{G}'$, set of bids $\{\phi_i\}_{i \in \mathcal{M}'}$, previous channel allocation $\mathbf{x}(t-1)$, continuous timeslots without preemption $\{\tau_i\}_{i \in \mathcal{M}'}$

Output: Channel allocation $\mathbf{x}(t)$, payment vector $\mathbf{p}(t)$

Compute optimal channel allocation $\mathbf{x}$

foreach $i \in \mathcal{M}'$ **do**

  Compute optimal channel allocation according to (8) $\mathbf{z}$ on $\mathcal{M}' \setminus \{i\}$

  $p'_i := \sum_{j \neq i \in \mathcal{M}'} \phi_j z_i(t) - \sum_{j \neq i \in \mathcal{M}'} \phi_j x_j(t)$

  $p_i(t) := \phi_i^{-1}\left(\frac{p'_i}{(1+\delta)^{\tau_i/T}}\right)$

Return $(\mathbf{x}(t), \mathbf{p}(t))$

Subject To:

$$\sum_{k \in setK} x_i^k(t) \leq 1 \qquad \forall i \in \mathcal{M}$$
$$\sum_{j \in setN_i} x_j^k(t) + x_i^k \leq 1 \quad \forall k \in \mathcal{K}, \forall i \in \mathcal{M}$$
$$x_i^k \in \{0, 1\} \quad \forall i \in \mathcal{M}, \forall k \in \mathcal{K}$$

The first constraint ensures that an agent is not assigned more than one channel at a time. The second constraint ensures that the assignment is interference free – agents that spatially interfere should not be assigned the same channel. Once a channel assignment is computed with the integer program in (8), we can compute a strategyproof payment mechanism by employing the VCG scheme.

B. The VCG Payment Scheme

Given a set of bids $\{\phi_i\}_{i \in \mathcal{M}'}$, the VCG mechanism chooses a channel allocation $\mathbf{x}^*$ that is optimal, and for each bidder i, computes a payment p'_i given as

$$p'_i = \sum_{j \neq i \in \mathcal{M}'} \phi_j z_i^* - \sum_{j \neq i \in \mathcal{M}'} \phi_j x_j^*, \tag{9}$$

where $\mathbf{z}^*$ is the optimal channel allocation obtained by setting $\phi_i = 0$ in the integer program in (8). See for example Nisan and Ronen [10] for more details on the VCG mechanism. Since the above payment is computed using the artificially inflated virtual valuations instead of the true values, the actual payment needs to be (i) scaled with respect to the inflation, and (ii) mapped back from virtual valuation space, as follows

$$p_i = \phi_i^{-1}\left(\frac{p'_i}{(1+\delta)^{\tau_i/T}}\right). \tag{10}$$

where ϕ_i^{-1} is the inverse virtual valuation function of i.

C. The Optimal Mechanism and Analysis

Algorithm 2 shows our mechanism in its entirety. The algorithm computes the optimal channel allocation that maximizes the valuations of bidders that are assigned a channel, while minimizing the cost of preemption. It then uses the VCG scheme to compute the payments specifically for timeslot t, which will be used to update the payment schedule $\mathbf{s}$ in Algorithm 1. In the following we show that Algorithm 1 is

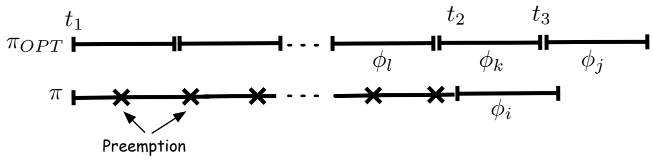

Fig. 2. Worst case preemption sequence used in proof for Theorem 3.

strategyproof when Algorithm 2 is used as the input monotone mechanism $\mathcal{A}$.

Theorem 2. The online auction framework of Algorithm 1, together with Algorithm 2, and the pricing scheme of (7), is strategyproof.

Proof: We need to show that the resulting mechanism is truthful for each bidder i in two dimensions:- the valuation v_i, as well as the temporal parameters a_i and d_i. This is achieved by showing that the allocation rule is monotone, and that the final price charged is the minimum payment required. Let θ_i and θ_i' be i's submitted bid when bidding truthfully and lying, respectively. We begin by showing that the allocation is monotone. If i bids truthfully and loses, either i is never assigned a channel, or is preempted one or more times between a_i and d_i. This means that i can never win by submitting a bid $v_i' < v_i$. So we must have $v_i = v_i'$. Furthermore, observe that the allocation decision is always independent of d_i. Therefore, we must have $d_i' = d_i$, since departing earlier can only harm i's chances of being allocated for T timeslots.

Now assume that she submits her bid at time $a_i' > a_i$, and ends up being allocated T continuous timeslots. If i was never allocated a channel at time $a_i \leq t < a_i'$ when bidding truthfully, we end up getting the same sequence of allocations as before, hence i will not be assigned a channel, which is a contradiction. On the other hand, if i was assigned a channel at time $a_i \leq t < a_i'$ in the truthful case, then there is now possibly a new sequence of allocations starting from a_i'. However, between a_i' and $d_i' = d_i$, i is preempted for the first time at some timeslot t' when bidding truthfully. Since delaying i's bid can only lower her virtual valuation for subsequent timeslots, i will once again be preempted at time t'. Therefore, there is no incentive to set $a_i' < t'$. Since we can recursively apply the previous argument starting from $a_i' \geq t'$, there is never an incentive for i to delay her bid.

Finally, observe that, by definition of the VCG scheme, the payment computed for each timeslot is the minimum payment required to be allocated a channel during that timeslot. Furthermore, the payments are independent from previous allocations, and the final price charged to every winning agent i in (7) is the minimum price chosen from all the valid timeslots. Therefore, misreporting $a_i' > a_i$ or $d_i' < d_i$ can only lead to an increase in the final price charged to i. We conclude that this auction mechanism is strategyproof. ■

The auction framework we have described in this section obtains a constant fraction of the optimal revenue even in the worse case, which is the focus of the next theorem. The crux of the following result is due to Hajiaghayi *et al.* [12]. We restate their proof in the context of an online spectrum auctions next.

Theorem 3. [Hajiaghayi *et al.*, 2005] The online auction framework of Algorithm 1 guarantees at least a $\frac{1}{5}$-fraction of the optimal revenue as computed by the offline adversary.

Proof: Without loss of generality, we can assume that the offline adversary is able to schedule perfectly, and hence no bids are preempted. On the other hand, bids in Algorithm 1 may be preempted. For every bid ϕ_j that was preempted at time t, it must be the case that there was another bid ϕ_i that was assigned j's channel instead, and furthermore, it must be the case that $\phi_i \geq (1+\delta)^{\tau_j}\phi_j$. The bid ϕ_j in turn may have preempted some other bid ϕ_k, and hence there may be a sequence of preemptions. Now pick one such sequence, call it π, and assume that π begins at timeslot t_1. Let ϕ_i be the last bid in π, which by definition does not get preempted, and let t_2 be the first of the T timeslots for which ϕ_i is assigned a channel. Between time t_1 and $t_2 + T$, there could be a sequence of non-preempted channel allocations by the optimal algorithm. Let us call this sequence π_{OPT}. Both sequences are shown in Figure 2. Let ϕ_k be the final bid in this sequence, and let t_3 be the time when this allocation takes place. Clearly, $t_3 \leq t_2 + T - 1$. Furthermore, by the monotonicity of the allocation chosen by Algorithm 2, we must have $\phi_k \leq (1+\delta)^{(T-1)/T}\phi_i$. Let ϕ_l be next highest bid in π_{OPT}, then $\phi_l \leq (1+\delta)^{-1/T}$ since ϕ_k and ϕ_l must be scheduled at least T timeslots apart. Proceeding in this fashion, we can construct the cost of the sequence of bids in π_{OPT} as

$$\begin{aligned} C &= (1+\delta)^{\frac{T-1}{T}}\phi_i + (1+\delta)^{\frac{-1}{T}}\phi_i + (1+\delta)^{\frac{1-T}{T}}\phi_i \ldots \\ &= \phi_i \left((1+\delta)^{1-\frac{1}{T}} + (1+\delta)^{\frac{-1}{T}} + (1+\delta)^{\frac{1}{T}-1} \ldots \right) \\ &= \frac{1}{(1+\delta)^{1/T}}\phi_i \left((1+\delta)^1 + 1 + (1+\delta)^{-1}(1+\delta)^{-2} \ldots \right) \\ &\leq \phi_i \left((1+\delta)^1 + 1 + (1+\delta)^{-1} + (1+\delta)^{-2} \ldots \right). \end{aligned}$$

Agents that leave the auction may return to bid again. When the number of users is large, the sequence above can go to infinity. The revenue approximation guarantee is then given by

$$\frac{\phi_i + C}{\phi_i} = 1 + (1+\delta)^1 + 1 + (1+\delta)^{-1} + (1+\delta)^{-2} \ldots$$

Setting $\delta = 1$ gives us

$$1 + 2 + 1 + (2)^{-1} + (2)^{-2} \ldots \leq 5,$$

which is the approximation ratio stated in the theorem. ■

Due to bid inflation, we are bounding the cost of preemption, and ensuring that preempting bids are always within a $O(1)$ factor of this cost. In fact, there is not much room for improving Theorem 3. It can be verified that different choices of δ offer only a fractional improvement. This bound is tight when the sequence of preemptions goes to infinity.

VI. A Greedy Approach to Channel Allocation

Since computing the optimal channel assignment is NP-Hard, the mechanism in Section V may be impractical for large problem instances. For the auction to be a truly scalable solution it is desirable to have an allocation mechanism that runs in polynomial time.

A. The Greedy Allocation Scheme and Channel Ranking

Given that the conflict graph can be modeled as a bounded disk graph, it is tempting to employ one of the many

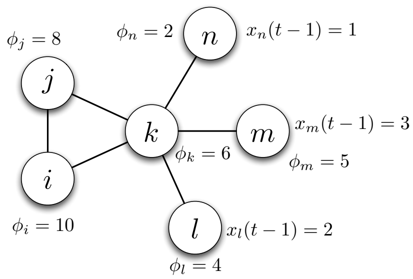

Fig. 3. Greedily allocating channels without considering the previous allocation causes unnecessary preemption.

polynomial-time approximation schemes (PTAS) for computing the maximum weighted independent set (MWIS) as a basis for the channel assignment algorithm. Unfortunately, both the readily available PTAS algorithms known for this problem, by Nierberg *et al.* [21] and Erlebach *et al.* [22], are not inherently monotone.

The simplest approximation algorithm that is also monotone is one based on the greedy approach. Consider the following simple allocation mechanism: (i) sort bidders in descending order, (ii) starting from the highest bidder, assign an available channel, (iii) if no channels are available, reject the bid and repeat from the previous step. This algorithm is clearly monotone, since reducing the bid can never cause a bidder to be considered for allocation earlier.

However, the allocation scheme described does not take into account the previous channel assignment. This can clearly cause unnecessary preemption. For an agent who has already been allocated a channel, the greedy scheme should try to assign the same channel, subject to availability. For agents being allocated for the first time, it is possible to create a *ranking* on the set of channels when considering them for assignment. For each agent i, channels should be ranked in the following order (from highest to lowest rank): (i) the channel assigned in the previous timeslot, if there is one, (ii) channels that have *not* been assigned to some agent $j \in \mathcal{N}_i$, and (iii) channels that have been assigned to some $j \in \mathcal{N}_i$, ranked in ascending order of their bids ϕ_j.

While this minimizes preemption, it still does not guarantee a good assignment. In Figure 3, there are three channels available, with the previous allocation shown as $\mathbf{x}(t-1)$. Assume that agents i, j and k only submit their bid starting at time t. The naive scheme would first assign channels 1 and 2 to agents i and j respectively, and then assign channel 3 to agent k, thus preempting m. The value of this solution is the sum of the assigned bids, which in this case is 30. We can improve this solution (to a solution with value 33) if we had assigned channel 1 to agent k instead thereby preempting the lowest bidder n instead. It is tempting to consider an improvement by first assigning channels, and then attempting to relabel them so as to minimize the cost of preemption. Unfortunately, this solution can fail to be monotonic, since the relabeling itself is a function of the previous allocation. The greedy algorithm has the advantage of being monotonic, and we will show that it still allows us to achieve a constant fraction of the optimal revenue.

B. Computing Truthful Payments Efficiently

As previously shown, the greedy algorithm leads to sub-optimal channel allocations. This precludes the use of the VCG payment scheme, since it is well known that VCG is not truthful when used together with an approximately optimal channel assignment [14]. Instead, we will tailor a truthful payment scheme specifically for our greedy allocation algorithm. Recall from Lemma 1, for any monotone algorithm, the payment for the winning bidder must be the minimum bid required to guarantee winning. One can in fact easily compute this in the following way: For each bidder i who has been allocated a channel, we set its bid $\phi_i = 0$, keep all other bids the same, and initialize $\mathcal{K}_i$ to be the set of available channels to i. We then simulate the algorithm again, while updating $\mathcal{K}_i$ when one of i's neighbors is allocated. The first agent j that is assigned a channel that empties $\mathcal{K}_i$ is i's *threshold agent*. Clearly, the threshold bid ϕ_j is the minimum bid required for i to win, and hence i should be charged this bid.

In fact, we will show an improved method for computing the payments directly during the allocation phase. Let us denote an agent i as being *saturated* if all channels have been assigned to the set consisting of i and $\mathcal{N}_i$. Now, observe that if $\phi(j)$ is the threshold bid for agent i, then j will not be assigned a channel. This means that whenever we are unable to assign a channel to some agent j, then it means that j is a threshold agent for another agent i that has *already been assigned a channel*. Moreover, if j is a threshold agent for i, then j is not only a neighbour of i, but it too must also saturated. Using these observations, it is now possible to compute the payments for agents during the allocation phase itself.

C. The Greedy Algorithm and Analysis

Algorithm 3 shows the greedy channel allocation scheme. The function

$$rank(i, \mathcal{N}_i, \mathcal{K}, \mathbf{x}(t), \mathbf{x}(t-1))$$

returns the first channel in the set of available channels, ranked in the manner described earlier so as to reduce preemptions. If no channels are available, it returns 0. The algorithm then proceeds to assign channels in a greedy manner. Whenever it finds that an agent i has been saturated, it computes a set of candidate agents, $\mathcal{C}$, for which i may be a threshold agent. It chooses the agent k in this set with the minimum bids, and then computes the payment for k based on ϕ_i.

Algorithm 3 is monotone by construction, and since the payments for winning agents are the minimum required bid to guarantee winning, the algorithm is clearly truthful. We state the next theorem without proof, since the reasoning is similar to the proof for Theorem 2, now that we have established that strategyproofness of Algorithm 3.

Theorem 4. The online auction framework of Algorithm 1, together with Algorithm 3, and the pricing scheme of (7), is strategyproof.

The performance bound of the greedy mechanism relies on the following geometric properties of disk graphs.

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

Algorithm 3: The greedy channel allocation and payment mechanism.

Input: Induced conflict graph $\mathcal{G}'$, set of bids $\{\phi_i\}_{i\in\mathcal{M}'}$, previous channel allocation $\mathbf{x}(t-1)$, continuous timeslots without preemption $\{\tau_i\}_{i\in\mathcal{M}'}$

Output: Channel allocation $\mathbf{x}(t)$, payment vector $\mathbf{p}(t)$

foreach $i \in \mathcal{M}'$ **do**
- $x_i(t) := 0$
- $p_i := 0$
- $done(i) :=$ false
- $saturated(i) :=$ false

Let $\mathcal{B} := \{\phi(i)\}_{i\in\mathcal{M}'}$

while $\mathcal{B} \neq \emptyset$ **do**
- $i := \arg\max_i\{\mathcal{B}\}$
- $\mathcal{B} := \mathcal{B} \setminus \{i\}$
- **if** $saturated(i)$ = *true* **then**
  - $\mathcal{T} := \{\phi_j | j \in \mathcal{N}_i \wedge done(j) = \text{false} \wedge x_j(t) > 0\}$
  - **foreach** $k \in \mathcal{T}$ **do**
    - $p'_k := \phi_i$
    - $p_k(t) = \phi_k^{-1}\left(\frac{p'_k}{(1+\delta)^{\tau_k/T}}\right)$
    - $done(k) :=$ true
- **else**
  - $x_i(t) := rank(i, \mathcal{N}_i, \mathcal{K}, \mathbf{x}(t), \mathbf{x}(t-1))$
  - **foreach** $j \in \mathcal{N}_i \cup \{i\}$ **do**
    - **if** *all* $\mathcal{K}$ *channels have been assigned to* $\{j\} \cup \mathcal{N}_j$ **then**
      - $saturated(j) :=$ true

Return $(\mathbf{x}(t), \mathbf{p}(t))$

This lemma can be derived by trying to fit as many agents with transmission radius R_{min} around an agent i with radius R_{max}, such that these agents do not interfere with one another, but interfere with i. In a greedy assignment scheme, the potential loss of of bids when allocating to i is therefore bounded by the constant $\lfloor\frac{2\pi}{\arcsin(\frac{1}{2\Delta}+1)}\rfloor - 1$, since all these neighbouring agents who could have been allocated instead must have bids at most ϕ_i. This observation, together with Theorem 3, leads to the following theorem.

Theorem 5. Let $C = 5\lfloor\frac{2\pi}{\arcsin(\frac{1}{2\Delta}+1)}\rfloor - 1$. Then Algorithm 3 used together with the Algorithm 1 guarantees at least a $\frac{1}{C}$-fraction of the optimal revenue.

While the formal proof is omitted due to space constraints, we note that the theorem follows directly from prior results. In summary, the greedy scheme achieves a constant revenue guarantee with respect to the optimal revenue of the offline adversary. For large-scale online auctions, the greedy algorithm, facilitated by our tailored truthful payment scheme, offers the advantage of a solution that is both easy to implement and fast.

VII. Conclusion

In this paper, we designed a revenue-maximizing online auction framework for dynamic secondary spectrum access. Our auction is provably strategyproof in two dimensions – the valuations of the bidder, as well as the timing of the bids. Since online allocations with irrevocable decisions can lead to an arbitrary loss of revenue, we introduce preemption into our system. Artificially inflating bids in the temporal dimension allows us to control the amount of preemption that takes place. We prove that this leads to a constant approximation mechanism with respect to the optimal revenue obtained by the offline adversary. This holds for both optimal channel allocation, as well as our tailored, truthful greedy scheme.

There are a number of interesting directions for future research. Our mechanism is only strategyproof when bidders act individually, and we would like to extend this to hold for the case of colluding bidders. We also only considered the case of binary demand, and it will be interesting to relax these assumptions when extending our framework. We intend to explore these ideas in future work.

References

[1] P. Crampton, "Spectrum auctions," in *Handbook of Telecommunications Economics*. Elsevier Science B.V., 2002, pp. 605–639.

[2] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *Proc. ACM MobiHoc*, 2009.

[3] X. Zhou, S.Gandhi, S.Suri, and H. Zheng, "eBay in the Sky: Strategyproof wireless spectrum auctions," in *Proc. ACM MOBICOM*, 2008.

[4] M. McHenry, P. Tenhula, D. McCloskey, D. Roberson, and C. Hood, "Chicago spectrum occupancy measurements analysis and a long-term studies proposal," in *Proc. ACM TAPAS*, 2006.

[5] S. Bridge, http://www.spectrumbridge.com/.

[6] L. Deek, X. Zhou, K. Almeroth, and H. Zheng, "To preempt or not: Tackling bid and time-based cheating in online spectrum auctions," in *Proc. IEEE INFOCOM*, 2011.

[7] A. Gopinathan and Z. Li, "A prior-free revenue maximizing auction for secondary spectrum access," in *Proc. IEEE INFOCOM*, 2011.

[8] A. Gopinathan, Z. Li, and C. Wu, "Strategyproof auctions for balancing social welfare and fairness in secondary spectrum markets," in *Proc. IEEE INFOCOM*, 2011.

[9] N. Nisan, T. Roughgarden, E. Tardos, and V. V. (Eds.), *Algorithmic Game Theory*. Cambridge University Press, 2007.

[10] N. Nisan and A. Ronen, "Algorithmic mechanism design (extended abstract)," in *Proc. ACM STOC*, 1999.

[11] P. Xu, S. Wang, and X. Li, "Salsa: Strategyproof online spectrum admissions for wireless networks," *IEEE Trans. on Comp.*, 2010.

[12] M. Hajiaghayi, "Online auctions with re-usable goods," in *Proc. ACM EC*, 2005.

[13] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-completeness*. W.H. Freeman and Company, 1979.

[14] N. Nisan and A. Ronen, "Computationally feasible VCG mechanisms," *Journal of Artificial Intelligence Research*, vol. 29, pp. 19–47, 2007.

[15] V. Krishna, *Auction Theory*. Academic Press, 2002.

[16] P. Klemperer, *Auctions: Theory and Practice*. Princeton University Press, 2004.

[17] R. Myerson, "Optimal auction design," *Mathematics of Operations Research*, pp. 58–73, 1981.

[18] 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.

[19] A. Borodin and R. El-Yaniv, *Online computation and competitive analysis*. Cambridge University Press, 1998.

[20] M. Hajiaghayi, R. Kleinberg, and D. Parkes, "Adaptive limited-supply online auctions," in *Proc. ACM EC*, 2004.

[21] T. Nieberg, J. Hurink, and W. Kern, "A robust ptas for maximum weight independent sets in unit disk graphs," in *Graph-Theoretic Concepts in Computer Science*. Springer, 2005, pp. 214–221.

[22] T. Erlebach, K. Jansen, and E. Seidel, "Polynomial-time approximation schemes for geometric graphs," in *Proc. ACM-SIAM SODA*, 2001.