A Prior-Free Revenue Maximizing Auction For Secondary Spectrum Access

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

***Abstract*—Dynamic spectrum allocation has proven promising for mitigating the spectrum scarcity problem. In this model, primary users lease chunks of under-utilized spectrum to secondary users, on a short-term basis. Primary users may need financial motivations to share spectrum, since they assume costs in obtaining spectrum licenses. Auctions are a natural revenue generating mechanism to apply. Recent design on spectrum auctions make the strong assumption that the primary user knows the probability distribution of user valuations. We study revenue-maximizing spectrum auctions in the more realistic *prior-free* setting, when information on user valuations is unavailable. A two-phase auction framework is constructed. In phase one, we design a strategyproof mechanism that computes a subset of users with an interference-free spectrum allocation, such that the potential revenue in the second phase is maximized. A tailored payment scheme ensures truthful bidding at this stage. The selected users then participate in phase two, where we design a randomized competitive auction and prove its strategyproofness through the argument of bid independence. Our solution applies iterative bidder partitioning on judiciously selected bidder subsets. Employing probabilistic techniques, we prove that our auction generates a revenue that is at least $\frac{1}{3}$ of the optimal revenue, improving the best known ratio of $\frac{1}{4}$ proven for similar settings.**

I. Introduction

Wireless spectrum is a scarce commodity, whose usage is subject to regulations and policies drawn up by governmental institutions such as the Federal Communications Commission (FCC) in the United States. The outcome is a *static* model of spectrum allocation, where large wireless operators compete through auctions [3], [8], [12]. Winners in such auctions are granted long-term licences, and are known as *primary spectrum users*. Recent studies suggest that such static allocation is rather inefficient [1], [15]. Spectrum utilization by primary users varies drastically, both geographically and temporally, with large areas and time periods without traffic. Unlicensed spectrum users, or *secondary users*, are faced with an artificial spectrum shortage [1], [21]. Recently, *dynamic spectrum allocation* [1], [10], [11], [21] has been proposed as a possible solution to this problem.

In the dynamic spectrum allocation model, a *spectrum broker* [22], [23] periodically pools unused portions of spectrum together. Separate chunks of this spectrum is then leased to secondary users, on a short-term basis [11], [23]. Recent studies have covered the design of fast spectrum allocation algorithms, and the feasibility and implementation issues of such a model [4], [22], [23]. However, a dynamic spectrum allocation protocol must realistically feature financial incentives for the primary user, since they accrue substantial costs in acquiring the spectrum license and in infrastructure [11], [23]. It is therefore desirable to incorporate an appropriate *revenue generating* mechanism into dynamic spectrum allocation schemes. A classic method for raising revenue in markets with limited goods is *auctions*, studied extensively in the economic literature [12], [13]. In supply-limited markets, the valuations by bidders for goods can be used by the auctioneer to set a pricing scheme, *e.g.* for achieving highest revenue. Furthermore, a well designed auction takes into account the possibility of *strategic behaviour* of bidders, who may misreport their true valuation, in the hopes of paying less for receiving the item. A *strategyproof* auction is one in which bidders have no incentive to lie about their true valuation [18], [20].

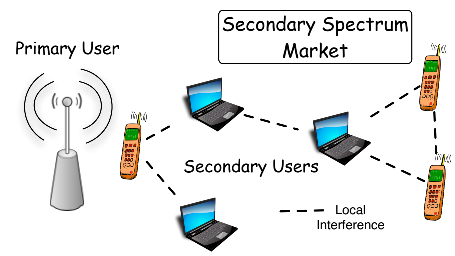

Fig. 1. Spectrum leasing from primary users to secondary users, through auctions. Goals of auction design: (i) ensure the spectrum allocation to secondary users is interference-free, and (ii) maximize revenue.

Fig. 1 shows a *secondary spectrum market*, where primary users lease spectrum to secondary users, using periodically held auctions. Such spectrum auctions have distinct *spatial* and *temporal* characteristics - secondary users are usually located in a geographically restricted area, and the spectrum lease is only valid for short durations [4], [11], [23], [25]. The proximity among secondary users implies that the auctioneer needs to ensure the spectrum allocation is *interference-free* [4], [11], [25]. Optimal interference-free allocation of spectrum is equivalent to the NP-hard problem of graph-colouring [5]. Efficient approximation algorithms for sub-optimal solutions are therefore necessary in practice.

However, employing a sub-optimal solution interferes with the desired strategyproof property of the auction [18], [19]. Zhou *et al.* [25] demonstrated that the well known Vickrey-Clarke-Groves auction [7], [24] fails to illicit truthful bidding in spectrum auctions with interference effects. They proceed to design a strategyproof one, focusing on maximizing the aggregate utility of secondary users. Nevertheless, dynamic spectrum auctions should be revenue-maximizing, in order to incentivize primary users to (temporarily) relinquish spectrum access. Jia *et al.* [11] propose an auction framework that is strategyproof and explicitly considers interference effects. Following the *de facto* standard of revenue-maximizing auction design in the economics [12], [13], [17], their mechanism requires the auctioneer to know probability distributions from which bidder valuations are drawn. This is a rather strong assumption, and is *unlikely* to hold in practical dynamic spectrum auctions. In such periodical, ephemeral auctions, a secondary user's valuation fluctuates over time, with little correlation. We are hence motivated to design a more pragmatic approach, where good revenue-generating properties do not depend on prior information on user valuations. To the best of our knowledge, this work is the first to consider spectrum auctions in the prior-free setting.

We target a revenue-maximizing spectrum auction in this *prior-free* setting [6], while explicitly taking into consideration the effects of interference. While our work is focused on dynamic spectrum auctions, our results and techniques may be of interest in other spectrum auction scenarios, where the key difficulty is the lack of information on bidder valuation. The proposed auction protocol is both *fast* and *strategyproof*: the algorithms run in polynomial time, and bidding truthfully is a provably dominant strategy for all users.

We consider interference-free allocation and revenue-maximization in two separate phases. In the first phase, we compute a feasible, interference-free channel assignment that potentially maximizes revenue based on the bids submitted by users. Only users with a channel assigned in the first phase remain in the auction. The second phase of the auction protocol consists of a strategyproof, randomized mechanism that computes a set of winners, who will receive the assigned channel at prices set to maximize revenue. While the second phase is strategyproof by design, the first phase is not intrinsically so. We resolve this problem by carefully tailoring a *payment scheme* to work with the assignment algorithm in the first phase, and prove its truthfulness. By treating the payments in the first phase as the *minimum price* a user must pay to win a channel in the second phase, we are able to ensure that the entire protocol is strategyproof.

The seminal work of Goldberg *et al.* [6] showed that no strategyproof auction can fully recover the optimal revenue in the prior-free setting [6]. Nonetheless, we show our auction achieves a constant, $\frac{1}{3}$-fraction of the optimal revenue, improving upon the random sampling auction introduced by Goldberg *et al.* [6] with guarantee of a $\frac{1}{4}$-fraction. A key idea introduced in our solution is to *partition* secondary users *iteratively*, and repeatedly extract revenue from a specific subset of them. The resulting mechanism is *truthful in expectation*, *i.e.*, the expected utility of bidders is maximized when bidding truthfully. The analysis of the revenue guarantee of this new auction is rather challenging. We employ probabilistic techniques to analyse the result of a related problem on the properties of a repeatedly partitioned set, from which we conclude that our auction asymptotically achieves a $\frac{1}{3}$-fraction of the optimal revenue in expectation. The techniques and results here are rather general, and may be of independent interest. In particular, we note that our prior-free revenue generating mechanism is widely applicable, and can be used for maximizing revenue in any auction setting where the supply of items is unlimited.

In the rest of the paper, we discuss related literature in Sec. II. The system model and some background on auction design are introduced in Sec. III. We present and analyze the randomized revenue-maximizing mechanism in Sec. IV, and study the interference-free channel assignment in a strategyproof manner in Sec. V. Sec. VI concludes the paper.

II. Previous Research

Recent studies on spectrum occupancy rates showed that statically allocated spectrum utilization varies drastically, both in the spatial and temporal dimensions [15]. In recent years, dynamic spectrum allocation has been suggested as a viable solution to efficiently utilize and share the available spectrum [1], [10], [21]. Dynamic spectrum access may be uncoordinated, with secondary users using cognitive radios for spectrum sensing [1], [10]. In contrast, the coordinated approach calls for a centralized entity, *e.g.*, a spectrum broker, to pool and manage unused spectrum, and to lease them to secondary users periodically [22], [23].

Revenue-maximization for dynamic spectrum auctions was considered by Gandhi *et al.* [4], who used a linear programming approach to model interference constraints. This work does not consider strategic user behavior, and assumes truthful bids for free. Strategic behaviour is considered by Sengupta and Chatterjee [22], who propose a knapsack based auction for dynamic spectrum allocation. However, they do not address the problem of interference. We target an auction protocol that explicitly finds an interference-free spectrum allocation, while maximizing revenue in a strategyproof fashion.

A *mechanism* is a protocol for implementing a desired *social choice* function. In this paper, we are primarily interested in mechanisms that are auctions. A classic treatment of auction theory can be found in the monograph of Krishna [13], while Klemperer [12] surveys practical and theoretical techniques used in modern auctions. The Vickrey-Clarke-Groves (VCG) mechanism [7], [24] is a powerful tool for designing strategyproof auctions, and is in some cases the only method for ensuring truthful behaviour [19]. However, VCG mechanisms lose the strategyproofness property when applied to suboptimal solutions, rendering it unsuitable for dealing with NP-Hard problems [18], [19].

In traditional economic theory, revenue-maximizing auctions, *a.k.a. optimal auctions*, are designed under the assumption that the auctioneer knows a probability distribution, from which bidder valuations are drawn. The influential work of Myerson [17] showed that applying the VCG mechanism using *virtual valuations* of bidders, computed using the valuations distribution, yields auctions that are both optimal and strategyproof. Relaxing the strong assumption on bidder valuation knowledge, Goldberg *et al.* [6] borrow techniques from the analysis of online algorithms to design optimal auctions. They show that no strategyproof and optimal auction exists in the paradigm of deterministic algorithms. Instead, they resort to randomized mechanisms, and show how a constant ratio of the optimal revenue can be recovered in expectation, within the context of selling digital goods with unlimited supply [6].

III. Preliminaries

We now introduce our system model in Sec.III-A and Sec.III-B, and describe some preliminaries in mechanism design in Sec.III-C.

A. System Model

This work concerns auction design in the *secondary spectrum market*, where the primary user is the *auctioneer*, and secondary users are bidding *agents*. The spectrum for auction is divided into a set $\mathcal{K}$ of *channels*, numbered $1, \ldots, K$. The *interference graph* $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ models the interference among the set of agents, denoted $\mathcal{S}$. Two agents $i, j \in \mathcal{S}$ interfere if they are assigned the same channel $k \in \mathcal{K}$, and $(i, j) \in \mathcal{E}$. The presence of edge (i, j) depends on the physical distance between i and j, and their transmitting power.

The auctioneer wishes to design an auction that allocates channels in an interference-free manner, and charges judiciously set prices to maximize total revenue. A *feasible channel assignment scheme* is a function $f : \mathcal{S} \to \mathcal{K} \cup \{0\}$, such that if $f(i) = f(j) \neq 0$, then $(i, j) \notin \mathcal{E}$. If $f(i) = 0$, then agent i has not been allocated a channel. A binary variable x_i to indicates whether agent i is assigned a channel: $x_i = 1$ if $f(i) \in \mathcal{K}$ and $x_i = 0$ otherwise.

B. Strategic Behaviour Model

Each agent $i \in \mathcal{S}$ has a *valuation* $v_i \in \mathbb{Q}^+$ for obtaining a channel, which is private information known only to itself. Here v_i can be interpreted in monetary terms, indicating the max amount agent i is willing to pay for a channel. We focus on the *prior-free* setting [6], in which no information on the distribution of the set of valuations $\{v_i\}$ is available. The revenue generating properties of the auction designed in this work hold for *any* distribution of agent valuations.

An agent i who 'wins', *i.e.* is allocated a channel, is charged a payment p_i by the auctioneer. Each agent i is assumed to be *selfish* and *rational* [13], [18], [20], trying to maximize its *utility* u_i:

$$u_i = \begin{cases} v_i - p_i & \text{if } i \text{ receives a channel} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

C. Strategyproof Mechanism Design

In the spectrum auction, each agent is required to disclose its valuation v_i at the beginning of the auction. The sole degree of freedom of an agent is the *value* of its bid. Agents may choose to mis-report their true valuation. We use b_i to denote agent i's bid. The vector of bids by all agents is $\mathbf{b}$, while the vector $\mathbf{b}_{-i}$ denotes the bids of all agents excluding i. Since the outcome of an auction is a function of all agents' bids, we can denote agent i's utility as $u_i(b_i, \mathbf{b}_{-i})$ to more precisely reflect this. An auction is *dominant-strategy truthful* if reporting the true valuation is the dominant strategy for an agent i, regardless of the bids of other agents. More formally, given a mechanism, for all $b_i \neq v_i$ and for any $\mathbf{b}_{-i}$, if the following always holds for all outcomes of the mechanism:

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

then the mechanism is dominant-strategy truthful, or *strategyproof*. Similarly, if Eq. (2) holds in expectation, then the mechanism is strategyproof in expectation.

The space of all possible revenue-maximizing mechanisms may be large. Yet we will restrict ourselves to mechanisms for which a *dominant-strategy implementation* [20] exists. Such mechanisms are characterized by outcomes in which every agent is playing her dominant strategy. This property helps us make concrete guarantees about the performance of the auction. We further focus on designing auctions that are dominant-strategy truthful. The *revelation principle* [17], [18] states that any mechanism implementable in dominant-strategies can be reduced to a dominant-strategy *truthful* mechanism. Besides truthfulness, we also target *individual rationality*, *i.e.*, $u_i \geq 0, \forall i$, and *no positive transfers*, *i.e.*, $p_i \geq 0, \forall i$ [20].

An important characterization of strategyproof mechanisms was given by Myerson [2], [17]:

Lemma 1. [Myerson, 1981] *Let $x_i(b_i)$ be the allocation function used in an auction for bidder i with bid b_i. A mechanism is strategyproof if and only if the following hold for a fixed $\mathbf{b}_{-i}$:*

- *$x_i(b_i)$ is monotonically non-decreasing in b_i*
- *Bidder i bidding b_i is charged $b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz$*

Lemma 1 indicates that once an allocation rule $x_i(.)$ for agent i is fixed in a strategyproof mechanism, it completely determines $i's$ payment. Conversely, given a fixed payment rule, the mechanism's allocation rule is determined accordingly. Lemma 1 thus provides two equivalent views to a truthful mechanism: (i) there exists a *critical bid* b'_i, such that i wins a channel iff 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 (i) when designing truthful mechanisms in Sec. V, while (ii) will be convenient when proving strategyproofness for the randomized mechanism proposed in Sec. IV.

IV. Phase 2: Revenue Maximization without Interference

Our strategy towards a revenue maximizing auction with interference-free property is to design a two-phase auction. In phase 1, we focus on computing a feasible (interference-free) channel assignment, such that the *potential* revenue from agents with winning potential, who advance into phase 2. A a randomized mechanism in phase 2 further computes the set of true winners, maximizing the revenue gleaned. We ensure that the strategyproof property holds individually for each each

Algorithm 1: A profit extraction procedure

Input: Target revenue R, set of agents $\mathcal{A}$
Output: Set of winners $\mathcal{W}$

1. Set $\mathcal{W} := \emptyset$
2. Let i be highest bidder in $\mathcal{A}$. Set $\mathcal{W} := \mathcal{W} \cup \{i\}$, set $\mathcal{A} := \mathcal{A} \setminus \{i\}$
3. Set price $p := \frac{R}{|\mathcal{W}|}$
4. If $v_i \geq p, \quad \forall i \in \mathcal{W}$, return $\mathcal{W}$. Otherwise repeat from step 2. If $\mathcal{A} = \emptyset$, return failure

phase. In Sec. V-C, we describe a method that guarantees the overall protocol is strategyproof over both phases.

We describe phase 1 of our protocol in Sec. V, and focus on phase 2 in this section instead. Within this section, we assume $|\mathcal{K}| = |\mathcal{S}|$, *i.e.*, there are as many channels as agents. Consequently, we can safely ignore interference effects (a feasible channel assignment can be found trivially), and focus on extracting the maximum revenue. In Sec. V, we will show how to compute an interference-free channel assignment in a strategyproof fashion, *without compromising* the revenue guarantees made in this section.

A. Optimal Pricing and Revenue Benchmark

We will design a revenue-maximizing mechanism that takes as input a set of agents for which a feasible assignment has already been found, denoted as $\mathcal{P}$. Hence, all agents in $\mathcal{P}$ may feasibly be declared winners in the auction, though such an outcome will most likely yield revenue that is arbitrarily far from optimal. We sort agents in $\mathcal{P}$ by their valuation, so that $v_i \geq v_j$ for all $i \leq j$. Since every agent can be feasibly assigned a channel, the *price* that maximizes revenue is

$$p = \arg\max_{v_i} \sum_{v_j \geq v_i} v_j \tag{3}$$

That is, the optimal price of Eq. (3) is the valuation v_i that maximizes iv_i, the total *single-price revenue*. The optimal single-price revenue, $R^{(1)}$, is:

$$R^{(1)} = \max_i iv_i \tag{4}$$

$R^{(1)}$ is non-convex, but a simple search through all v_i, $i = 1 \ldots |\mathcal{P}|$, finds the optimal revenue in polynomial time.

In the assumed prior-free setting, the auctioneer has no information on the distribution of agent valuations in $\mathcal{P}$. The first obstacle in such a setting is to design a mechanism that can 'guess' the optimal price to charge agents. Setting p either too high or too low may lose revenue. The lack of information on user valuations can be remedied if we use a subset of bids as a *sample*, to gain a good estimate on the optimal revenue maximizing price p. We can then offer channels at the price p, but only to agents whose bids were not sampled, thus discarding the potential revenue from this sampled subset of agents. Lemma 1 assures that since the price offered to each agent is not a function of its bid, this mechanism is indeed strategyproof.

The second obstacle in the prior-free setting is deciding a suitable revenue benchmark, to compare the revenue obtained by our mechanism with. The single-price revenue of Eq. (4) has the advantage that one can design strategyproof mechanisms to extract a constant amount of this revenue. Further, Immorlica *et al.* [9] have shown that this revenue is within a constant factor of any differential pricing mechanism privy to the distribution of user valuations. However, in the case that the maximum revenue $R^{(1)}$ of Eq. (4) is obtained by allocating a channel to only *one* bidder, it is impossible for any strategyproof auction to generate a constant fraction of $R^{(1)}$ [6]. This negative result can be easily deduced from Lemma 1. Since the prices charged to i in any truthful auction must be a function of $\mathbf{b}_{-i}$, and the optimal revenue in this case is b_i, one cannot hope to obtain a constant fraction of b_i. Consequently, we will adopt the previous convention in prior-free auctions and consider as our benchmark the revenue obtained from the single-price auction where *at least* two agents are allocated [6]. Formally, this is defined as:

$$R^{(2)} = \max_{i \geq 2} iv_i \tag{5}$$

Algorithm 2: Iterative Random Partitioning Mechanism

Input: Set of agents, $\mathcal{P}$
Output: Total revenue R

1. Initialize $R := 0$
2. Randomly partition bidders in $\mathcal{P}$ to two sets, $\mathcal{A}$ and $\mathcal{B}$
3. Compute optimal price and revenue for sets $\mathcal{A}$ and $\mathcal{B}$ using Eq. (3) and Eq. (4). Let $R_{\mathcal{A}}$ and $R_{\mathcal{B}}$ denote the revenue from sets $\mathcal{A}$ and $\mathcal{B}$ respectively
4. Run Algorithm 1 on $\mathcal{A}$ (resp. $\mathcal{B}$) using target revenue $R_{\mathcal{B}}$ (resp. $R_{\mathcal{A}}$)
5. Let R' be revenue gained from running Algorithm 1 successfully
6. Set $R := R + R'$
7. Let $\mathcal{P}'$ denote one of two sets $\mathcal{A}$ and $\mathcal{B}$ from which Algorithm 1 failed to raise any revenue
8. If $|\mathcal{P}'| > 1$, set $\mathcal{P} := \mathcal{P}'$. Repeat from step 2
9. Otherwise return R

B. A Revenue-Maximizing Randomized Auction

A key ingredient of our revenue-maximizing mechanism is a Cournot tattonnement [16] for *profit extraction*, shown in Algorithm 1. It takes as input a set of agents $\mathcal{A}$, and a *target revenue* R. It attempts to compute a subset of agents that are able to share R equally, through gradual elimination of agents that are unable to afford the offered price $\frac{R}{|\mathcal{A}|}$. First, the procedure is strategyproof, since the offer price is a function of a fixed R and the number of remaining agents in $\mathcal{A}$, and is independent of agents' bids. Second, if R^* is the optimal single-price revenue of Eq. (4) that can be extracted from the set of agents $\mathcal{A}$, then clearly the procedure will successfully *extract* R^* from $\mathcal{A}$ if and only if the target revenue $R \leq R^*$. Algorithm 1 is similar to the profit extraction method of Goldberg *et al.* but assumes offer price is descending in every round instead of ascending. We will see later that this subtle modification is crucial towards ensuring that the iterative partitioning auction we later design is strategyproof in expectation.

Armed with this profit extraction procedure, we are ready to describe our mechanism, shown in Algorithm 2. The mechanism begins by randomly *partitioning* the set of agents $\mathcal{P}$ into two sets, $\mathcal{A}$ and $\mathcal{B}$. It computes the optimal revenue for each set $R_{\mathcal{A}}$ and $R_{\mathcal{B}}$ based on the submitted bids of agents, using Eq. (4). The mechanism then attempts to extract $R_{\mathcal{A}}$ from $\mathcal{B}$, and vice versa, using the mechanism of Algorithm 1. Since agents in one partition is offered a price that depends on the optimal revenue computed on the *other* partition, the price is *bid independent* for all agents. From Lemma 1, the mechanism is strategyproof at this point. Since we must have either $R_{\mathcal{A}} \geq R_{\mathcal{B}}$ or $R_{\mathcal{B}} \geq R_{\mathcal{A}}$, we are guaranteed from the properties of Algorithm 1 that at least one partition will yield revenue, specifically the minimum of $(R_{\mathcal{A}}, R_{\mathcal{B}})$. If Algorithm 1 successfully extracts $R_{\mathcal{B}}$ ($R_{\mathcal{A}}$) from $\mathcal{A}$ ($\mathcal{B}$), then we assign channels to the winning bidders in this partition at the prices computed by the profit extraction procedure. We repeat the entire process, but considering only bidders from the *partition that did not yield any revenue* during the profit extraction phase. Later, we will see that our choice of agents that participate in the next round, together with the fact that Algorithm 1 offers in descending order, leads to a mechanism that is strategyproof in expectation.

C. Analysis of Revenue Guarantee

We next analyze the revenue guarantee of our mechanism. Let us denote by p^* the price that maximizes the revenue benchmark shown in Eq. (5), and assume that w is the number of agents with valuations at least p^*. Then, the optimal revenue available is $R^{(2)} = wp^*$. Our goal is to compute the fraction of this revenue obtained by Algorithm 2. Let $R(j)$ be the revenue generated by Algorithm 2 in round j of the mechanism, and let $w(j)$ be the number of agents participating in round j who have valuation *at least* p^*. Then, the portion of $R^{(2)}$ in round j, $R^{(2)}(j)$, is:

$$R^{(2)}(j) = p^* w(j) \tag{6}$$

Let $w_{\mathcal{A}}(j)$ and $w_{\mathcal{B}}(j)$ be random variables expressing the number of agents that end up in the sets $\mathcal{A}$ and $\mathcal{B}$ respectively, during iteration j, such that $w_{\mathcal{A}}(j)+w_{\mathcal{B}}(j) = w(j)$. We define the random variable $\alpha(j)$ as

$$\alpha(j) = \min\Big(\frac{w_{\mathcal{A}}(j)}{w(j)}, \frac{w_{\mathcal{B}}(j)}{w(j)}\Big) \tag{7}$$

Combining this with the properties of the profit extraction procedure shown in Algorithm 1, leads to the following lemma:

Lemma 2. *Given $w(j)$, Algorithm 2 generates at least a α-fraction of $R^{(2)}(j)$ in expectation,* i.e. $\frac{E[R(j)]}{R^{(2)}(j)} = E[\alpha]$.

Proof: Let $R_{\mathcal{A}}(j)$ and $R_{\mathcal{B}}(j)$ be the optimal revenue computed for sets $\mathcal{A}$ and $\mathcal{B}$ respectively during iteration j. From the profit extraction procedure of Algorithm 1, we know that we are guaranteed the *lesser* of the two quantities, *i.e.* $R(j) = \min(R_{\mathcal{A}}(j), R_{\mathcal{B}}(j))$. Also, it must be that $R_{\mathcal{A}}(j) \geq w_{\mathcal{A}}(j)p^*$ and $R_{\mathcal{B}}(j) \geq w_{\mathcal{B}}(j)p^*$, and we get

$$\begin{aligned}\frac{E[R(j)]}{R^{(2)}(j)} &= \frac{E\big[\min(R_{\mathcal{A}}(j), R_{\mathcal{B}}(j))\big]}{R^{(2)}(j)} \\ &= \frac{E\big[\min(w_{\mathcal{A}}(j)p^*, w_{\mathcal{B}}(j)p^*)\big]}{w(j)p^*} \\ &= \frac{E\big[\min(w_{\mathcal{A}}, w_{\mathcal{B}}(j))\big]}{w(j)} = E[\alpha(j)]\end{aligned}$$

■

Essentially, Lemma 2 allows us to reframe the revenue guarantee in any iteration of the mechanism, in terms of the min number of agents over both set partitions $\mathcal{A}$ and $\mathcal{B}$ who have valuations that are now less than p^*, the revenue-maximizing price . We next develop a theorem that allows us to state a concrete value for $\alpha(j)$.

Theorem 1. *Let n be a set of balls, with $k \geq 2$ black balls and $n-k$ white balls. In a random partition of this set of balls, denote by X the random variable denoting the minimum number of black balls in any one partition. Then $E[X] \geq \frac{k}{4}$.*

Proof: We focus on the case that k is odd – a similar proof applies when k is even. There are $\binom{n}{n/2}$ ways of choosing balls for any one partition. Let j be the number of black balls chosen, then j is at most $\lfloor\frac{k}{2}\rfloor$ for the event we are interested in. Given that the minimum number of black balls occurring in either partition is a disjoint event (since k is odd), we get

$$E[X] = \sum_{j=1}^{\lfloor\frac{k}{2}\rfloor} \frac{2j\binom{k}{j}\binom{n-k}{n/2-j}}{\binom{n}{n/2}}$$

Since $j \leq \lfloor\frac{k}{2}\rfloor$, we claim that

$$\frac{2j\binom{k}{j}\binom{n-k}{n/2-j}}{\binom{n}{n/2}} \geq \frac{j}{2^{k-1}}\binom{k}{j}$$

The proof for the above can be found in the appendix. Next, we compute the sum of the above term for $j = 1 \ldots \lfloor\frac{k}{2}\rfloor$:

$$\sum_{j=1}^{\lfloor\frac{k}{2}\rfloor} \frac{1}{2^{k-1}} j\binom{k}{j} \geq \sum_{j=1}^{\lfloor\frac{k}{2}\rfloor} \frac{1}{2^{k-1}}\frac{k+1}{4}\binom{k}{j}$$

$$= \frac{k+1}{(4)(2^{k-1})}\Big[\frac{2^k}{2} - \binom{k}{0}\Big] = \frac{k+1}{4} - \frac{1}{4}\frac{k+1}{2^{k-1}} \geq \frac{k}{4}$$

The first inequality holds since the sum of the coefficients j in front of the binomial term is lower bounded by the average of these coefficients. ■

We can think of the entire set of balls as agents participating in an iteration of Algorithm 2, with black balls representing agents with valuations at least p^*. Theorem 1 then tells us that the ratio of the min number of black balls to all the balls in any one partition is at least $\frac{1}{4}$ in expectation. This, together with Lemma 2, implies that in expectation we make at least a $\frac{1}{4}$-fraction of the optimal revenue $R^{(2)}$ in each iteration of Algorithm 2.

We are now ready to state our main theorems. The first shows that Algorithm 2 achieves in expectation a $\frac{1}{3}$-fraction of the optimal revenue asymptotically, while the second proves that our mechanism is strategyproof in expectation.

Theorem 2. The expected revenue generated by Algorithm 2 is given by $R^{(2)}(\frac{1}{3} - \frac{1}{12n^2})$, where $R^{(2)}$ is the optimal single-price revenue when agent valuations are known exactly, and $n = |\mathcal{P}|$.

Proof: We know from Lemma 2 that $\frac{E[R(j)]}{R^{(2)}(j)} = E[\alpha(j)]$. If we think of the black balls as the number of agents who contribute to the optimal revenue, Theorem 1 implies that $\alpha(j) \geq \frac{1}{4}$ for all j. If we obtained revenue $R(j)$ in round j, then this is the optimal revenue from the *losing set*, which gets to participate in round $j+1$. The expected revenue in round $j+1$, conditional on the revenue obtained in round j, is then at least a $\frac{1}{4}$-fraction of the latter:

$$E[R(j+1)|R(j)] \geq \frac{R(j)}{4}$$

Taking the expectation on both sides, we get

$$E\Big[E[R(j+1)|R(j)]\Big] \geq \frac{E[R(j)]}{4}$$

and applying the law of iterated expectations we get

$$E[R(j+1)] \geq \frac{E[R(j)]}{4}$$

Let us assume that the number of agents, $|\mathcal{P}| = n$ is a power of 2. This is without loss of generality, since if not, the auctioneer can add dummy bidders with valuation 0 each. The number of iterations in the mechanism is therefore $\log_2 n$. Since in expectation we make a $\frac{1}{4}$-fraction of the revenue of iteration j in iteration $j+1$, the expected revenue of this mechanism, $E[R]$, is:

$$E[R] \geq \sum_{j=1}^{\log_2 n} \frac{R^{(2)}}{4^j} = \frac{R^{(2)}}{4} \frac{1-(1/4)^{\log_2 n - 1}}{1 - 1/4}$$

$$= \frac{R^{(2)}}{3}(1 - 2^{-2\log_2 n}2^{-2}) = R^{(2)}(\frac{1}{3} - \frac{1}{12n^2})$$

The first inequality is based on the linearity of expectations. ■

Theorem 3. Algorithm 2 is strategyproof in expectation.

Proof: In the first round, agents are partitioned randomly, and offered prices according to Algorithm 1 which is independent of its bid. For subsequent rounds, we need to show that the offer price stays the same in expectation. Using a similar argument to the proof of theorem 2, we know that the expected profit in round i is $\frac{k_i}{4}p^*$, where p^* is the optimal single-price revenue, and k_i is the number of agents in round i with valuations greater than p^*. But this means that Algorithm 1 attempts to extract $\frac{k_i}{4}p^*$ from $\frac{3k_i}{4}$ agents that can afford p^* in expectation, which implies the expected price is not lower than p^*. Therefore, in expectation, each agent's best strategy is to bid truthfully. ■

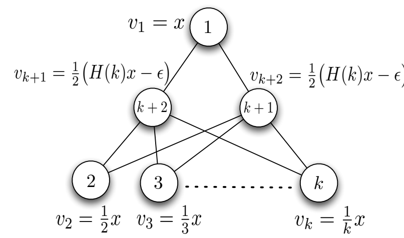

Fig. 2. An example topology where picking the set of agents which maximize the total agent valuations results in poor single-price revenue.

V. Phase 1: Strategyproof Interference-Free Channel Assignment

Algorithm 2 required as input a set of agents, all of whom can be assigned a channel feasibly. In practice, channels are a limited resource, and finding a feasible channel assignment is a computationally difficult problem, equivalent to graph colouring. In this section, we will design a mechanism that employs *any given* graph colouring algorithm for computing channel assignments in a *strategyproof* fashion, such that the potential revenue to be gained from the second phase of the protocol is not compromised.

A. Computing A Revenue-Maximizing Channel Assignment

Consider the case when all agents have the same valuation. Then finding a revenue-maximizing feasible assignment would be equivalent to maximizing the *number* of agents that can be feasibly allocated a channel. Given an interference graph and a fixed number of available channels, this reduces to the well known Maximum-K-Colourable Induced Subgraph problem (`MAX-K-CIS`) [5]. More formally, given a graph $\mathcal{G}$ and a set of colours $\mathcal{K}$, the `MAX-K-CIS` problem is that of computing a subgraph $\mathcal{H}$ of $\mathcal{G}$ that can be coloured with at most $|\mathcal{K}|$ colours, such that the number of vertices in $\mathcal{H}$ is *maximized*.

The `MAX-K-CIS` problem is NP-Hard, the special case of $K = 1$ reduces to the well known Maximum Independent Set problem, which is also NP-Hard [5]. Nevertheless, polynomial-time algorithms with good approximation factors exist for special cases of graphs such as unit disk graphs [14], which model the case when secondary users have uniform antennas. We will not take the approach of favouring any particular algorithm that solves `MAX-K-CIS`. Instead, a central design goal of the mechanism developed in this section is to ensure that it remains compatible with *any* algorithm, $\mathcal{A}$, that solves `MAX-K-CIS`. We require only that $\mathcal{A}$ is a *deterministic* algorithm. This design decision provides the auctioneer with the freedom to plug any (deterministic) algorithm $\mathcal{A}$ into our auction protocol. Thus the auctioneer can employ the algorithm $\mathcal{A}$ which provides the best performance given the properties of the relevant interference graph.

Of course, agents do not all share the same valuation for obtaining a channel. In this case, the naive approach would be to consider solving the *weighted* version of `MAX-K-CIS`, where the goal is to find a K-colorable subgraph of $\mathcal{G}$ such that the *sum of agent valuations* in the subgraph is maximized. This appears to be a reasonable approach, since the sum of agent valuations provides an upper bound on the amount of revenue that can be extracted. However, such a solution surprisingly does not maximize the single-price revenue of Eq. (4), when interference effects are taken into consideration. We illustrate this with the simple example in Fig. 2. In this interference graph, there is only 1 channel available for assignment (so that the problem reduces to finding a maximum weighted independent set). If we seek an assignment that maximizes the sum of agent valuations, the optimal solution would assign the channel to agents $1 \dots k$, for a total valuation of $x + \frac{1}{2}x + \dots + \frac{1}{k}x = H(k)x$, where $H(k) = O(\ln k)$ is the harmonic number of k. However, the maximum single-price revenue of Eq. (4) that can be extracted from this set of agents is exactly x. If we pick agents $k+1$ and $k+2$ to be allocated the channel instead, the maximum single-price revenue is obtained as $H(k)x - \epsilon$, where $\epsilon > 0$ is a quantity arbitrarily close to 0. Clearly, choosing a solution that aims to maximize the total agent valuation may reduce the maximum single-price revenue available by a factor of $\frac{H(k)x-\epsilon}{x} = O\big(H(k)\big)$.

The profit extraction procedure in Algorithm 1 is designed to extract the maximum available single-price revenue. This, together with the example shown in Fig. 2, illustrates the need to find an assignment that explicitly maximizes the latter quantity. Observe that when a price p is fixed, we need to assign channels to as many agents as feasible whose valuations $v_i \geq p$, in order to maximize Eq. (4). That is, it is sufficient to solve the MAX-K-CIS problem on the interference graph induced on the subset of agents with $v_i \geq p$. If $\mathcal{W}$ is the set of agents successfully assigned a channel by some algorithm $\mathcal{A}$, then the maximum single-price revenue available from this set is given as $|\mathcal{W}|p$.

From Eq. (3), we know that the optimal price p^* can only take on finitely many values, specifically, a value v_i for some $i \in \mathcal{S}$. This immediately points to a simple recipe for computing an assignment that maximizes revenue. Repeatedly run $\mathcal{A}$ on the subgraph of $\mathcal{G}$ induced by agents with valuation at least p, for each *threshold price* $p = v_1 \dots v_n$, where $n = |\mathcal{S}|$. If $\mathcal{W}_p$ is the set of agents that can be feasibly allocated a channel by $\mathcal{A}$ when the threshold price is p, then the revenue at this price is simply $p|\mathcal{W}_p|$.

We detail the procedure for computing a revenue-maximizing assignment in Algorithm 3. The algorithm takes as input the interference graph $\mathcal{G}$, the set of channels $\mathcal{K}$, as well as an algorithm $\mathcal{A}$ for solving MAX-K-CIS. The algorithm operates on a *base set* of agents $\mathcal{T}$, where $\mathcal{T} = \mathcal{S}$ initially. In subsequent iterations, $\mathcal{T} = \mathcal{S} \setminus i$, for each i, *i.e.*, we also compute the optimal revenue available *without the participation* of each agent i. This allows us to measure the *externality* [20] imposed by each agent i on other agents. In Sec. V-B, we show how this information is essential in the design of a payment scheme to ensure strategyproofness is preserved during this assignment phase.

For each such base set of agents, a threshold price $p = v_i$ is fixed for some valuation v_i such that $i \in \mathcal{T}$, and a subgraph is induced on $\mathcal{G}$ consisting only of agents in the base set $\mathcal{T}$, whose valuations are at least p. The algorithm $\mathcal{A}$ is then used

Algorithm 3: Algorithm for finding an interference free channel assignment that maximizes revenue R

Input: Interference graph $\mathcal{G}$, an algorithm $\mathcal{A}$ for the Max-K-CIS problem, set of channels $\mathcal{K}$
Output: A channel assignment f^*, allocated agents $\mathcal{W}^*$, such that single-price revenue R^* is maximized

Initialize $\mathcal{O} := \emptyset$, $\mathcal{T} := \mathcal{S}$, $R_{MAX} := 0$;
for $k = 0 \dots |\mathcal{S}|$ **do**
  if $k > 0$ **then**
    $\mathcal{T} := \mathcal{S} \setminus \{k\}$;
  for $i \in \mathcal{T}$ **do**
    Initialize $p := b_i$, $R := 0$, $\mathcal{W} = \emptyset$;
    $\mathcal{U} = \{i \in \mathcal{T} | b_i \geq p\}$;
    $\mathcal{G}'$:= Subgraph of $\mathcal{G}$ induced by agents in $\mathcal{U}$;
    Run $\mathcal{A}$ on $\mathcal{G}'$ with $|\mathcal{K}|$ colours. Let $\mathbf{f}$ be the resulting assignment ;
    for $j \in \mathcal{U}$ **do**
      if $f(j) > 0$ **then**
        $R := R + p$;
        $\mathcal{W} := \mathcal{W} \cup \{j\}$;
    if $R > R_{MAX}$ **then**
      $\mathbf{f}^* = \mathbf{f}$;
    $\mathcal{O} := \mathcal{O} \cup \{(R, \mathcal{W}, \mathcal{T})\}$;
$(R^*, \mathcal{W}^*, \mathcal{T}^*) := \max_R (R, \mathcal{W}, \mathcal{T}) \in \mathcal{O}$;
Output $(R^*, \mathcal{W}^*, \mathcal{T}^*)$ and $\mathbf{f}^*$;

to compute a channel assignment $\mathbf{f}$, such that the number of agents that can be feasibly assigned K channels is maximized. Let $\mathcal{W}$ be the set of agents allocated a channel by $\mathcal{A}$. The outcome of each iteration is stored as a triple $(R, \mathcal{W}, \mathcal{T})$ with revenue R, set of allocated agents $\mathcal{W}$, and the base set of agents $\mathcal{T}$. At the end, the algorithm returns the outcome that achieves the highest revenue R as the final solution.

B. Ensuring Strategyproofness for the Channel Assignment

While Algorithm 3 is fairly straightforward, we have so far neglected the implications of strategic behaviour by agents. Algorithm 3 works only when agents bid truthfully, *i.e.* $b_i = v_i$ for all i. The next step is therefore to ensure strategyproofness. The key difficulty here is that the purpose of computing an assignment in this phase is to find a suitable set of agents that can participate in the revenue-maximizing mechanism described in Sec. IV. Hence, not all agents who can be feasibly assigned a channel will end up being allocated one at the end of both phases of the auction protocol. Our strategy for now will be to ensure that the assignment mechanism is independently strategyproof instead, *i.e.* if agents are indeed allocated channels according to Algorithm 3, then the mechanism is truthful. Later, we will show in Sec. V-C how to ensure strategyproofness over both phases of the auction.

Algorithm 3 can be made strategyproof with the design an of accompanying *payment scheme*. Each agent i will be charged a payment p_i, such that the dominant strategy for each agent is to bid truthfully when Algorithm 3 is used to compute channel allocations. We note that the classic VCG payment scheme [7], [24] cannot be used in this setting, since the metric to maximize is not the total valuation of agents. Instead, our payment scheme must be tailored specifically for use with Algorithm 3. Algorithm 4 shows our protocol

Algorithm 4: Strategyproof payment scheme for Algorithm 3

Input: Set of all agents $\mathcal{S}$, a channel assignment $\mathbf{f}^*$ and a collection of outcomes $\mathcal{O}$ as computed by Algorithm 3
Output: Strategyproof pricing scheme $\mathbf{p}$

for $i \in \mathcal{S}$ **do**
    $R_{-i} := \arg\max_R \{(R, \mathcal{W}, \mathcal{T}) \in \mathcal{O} \mid i \notin \mathcal{T}\}$;
    $(R'_i, \mathcal{W}'_i, \mathcal{T}'_i) := \min_R \{(R, \mathcal{W}, \mathcal{T}) \in \mathcal{O} | R \geq R_{-i} \wedge i \in \mathcal{W}\}$;
for $i = 1 \ldots \mathcal{S}$ **do**
    if $f(i) > 0$ **then**
        $p_i := \frac{R_{-i}}{|\mathcal{W}'_i|}$;
    else
        $p_i := 0$;
Output $\mathbf{p}$;

for computing agent payments. The algorithm takes as input the set of solutions $\mathcal{O}$ as computed by Algorithm 3. It then computes the quantity R_{-i} as:

$$R_{-i} = \arg\max_R \{(R, \mathcal{W}, \mathcal{T}) \in \mathcal{O} \mid i \notin \mathcal{T}\} \tag{8}$$

That is, R_{-i} is the maximum achievable revenue when agent i is not part of the base set $\mathcal{T}$. It is now clear why we computed the maximum revenue available without each agent i in Algorithm 3. We say an outcome $(R'_i, \mathcal{W}'_i, \mathcal{T}'_i) \in \mathcal{O}$ is a *critical outcome* for agent i if and only if the three conditions hold: (i) $i \in \mathcal{W}'_i$, (ii) $R'_i \geq R_{-i}$, and (iii) For any other outcome $(R''_i, \mathcal{W}''_i, \mathcal{T}''_i) \in \mathcal{O}$ if $R''_i \geq R_{-i}$ and $i \in \mathcal{W}''_i$, then $R''_i \geq R'_i$. Hence, the critical outcome of some agent i is one that achieves at least as much revenue as the highest revenue generated without agent i's participation, and further this revenue is the smallest of all such outcomes.

Next, we will prove that under the payment scheme computed by Algorithm 4, it is a dominant strategy for each agent to report its true valuation. The proof is built upon three lemmas. We begin with one that pertains to the properties of the solution computed by Algorithm 3.

Lemma 3. The assignment rule in Algorithm 3 for some agent i is monotonically non-decreasing in b_i, for a fixed $\mathbf{b}_{-i}$. That is, if agent i is allocated a channel when bidding b_i with all other bids fixed at $\mathbf{b}_{-i}$, then agent i will be assigned a channel when bidding any $b'_i \geq b_i$.

Proof: Let R and $\mathcal{W}$ be the revenue and winner set respectively when i bids b_i. Similarly, let R' and $\mathcal{W}'$ be the revenue and set of winning agents when i bids b'_i instead. By way of contradiction, assume that $i \notin \mathcal{W}'$. Since $\mathbf{b}_{-i}$ is fixed and $b'_i > b_i$, the solutions returned by any deterministic algorithm $\mathcal{A}$ for threshold prices $p < b_i$ must be the same for both bids, since the input set of agents $\mathcal{U}$ remains the same. Therefore, the threshold price p that yields revenue R' must be greater than b_i. However, since $i \notin \mathcal{W}'$, there is a feasible assignment for the set of agents in $\mathcal{W}'$, with threshold price $p > b_i$ yielding higher revenue than R, contradiction. ∎

Next, we prove that our payment scheme satisfies the individual rationality property:

Lemma 4. For all agents i, the prices p_i computed by Algorithm 4 is individually rational when agents bid truthfully, that is $u_i(v_i) \geq 0$.

Proof: If an agent i is not allocated a channel, then $p_i = 0$ and the lemma holds. If an agent is assigned a channel, then let $(R'_i, \mathcal{W}'_i, \mathcal{T}'_i)$ be the critical outcome for i. We must have $R^* \geq R'_i \geq R_{-i}$. Furthermore, since $i \in \mathcal{W}'$, we get $v_i \geq \frac{R'_i}{|\mathcal{W}'|} \geq \frac{R_{-i}}{|\mathcal{W}'|} = p_i$ ∎

Lemma 5. Assume that Algorithm 3 assigns a channel to agent i both when i bids its true valuation v_i, and when i bids $b_i \neq v_i$, with all other bids $\mathbf{b}_{-i}$ fixed. Let $(R'_i, \mathcal{W}'_i, \mathcal{T}'_i)$ and $(R''_i, \mathcal{W}''_i, \mathcal{T}''_i)$ be the critical outcomes for when i bids v_i and b_i respectively. Then $\mathcal{W}'_i = \mathcal{W}''_i$.

Proof: The threshold price p' used for obtaining the set $\mathcal{W}'_i$ must satisfy $p' \leq v_i$. First assume $b_i > v_i$. Then when i bids b_i, since all other bids are fixed, the critical outcome, the one with the *minimum* revenue $R''_i \geq R_{-i}$, will still have the same price threshold $p' \leq v_i < b_i$. Since $\mathcal{A}$ is deterministic, the same solution $\mathcal{W}'$ will be picked as a feasibly allocated set of agents when i bids b_i. When $b_i < v_i$, we have $b_i \geq p''$, where p'' is the price threshold used when picking $\mathcal{W}''$. However, when i bids $v_i > p''$, once again we have the same symmetric situation as the previous case, where the outcome with the minimum revenue $R'_i \geq R_{-i}$ must be an outcome with price threshold $p' = p''$. Since $\mathcal{A}$ is deterministic, the same solution $\mathcal{W}''$ is picked as well. ∎

Theorem 4. The payment scheme shown in Algorithm 4 used in conjunction with Algorithm 3 for allocating channels to agents is strategyproof.

Proof: Let v_i and b_i be the true and declared valuation of some agent i. Similarly, let $x(v_i)$ and $x(b_i)$ be binary variables denoting the assignment outcome (*i.e.* assigned a channel or not) of Algorithm 3 when agent i bids v_i and b_i respectively. We denote by $p(v_i)$ and $p(b_i)$ the resulting payment charged to agent i when bidding either v_i or b_i. We will prove that under the pricing scheme of Algorithm 3, it is a dominant strategy for each agent to declare its valuation truthfully, by comparing the utility of agent i on a case-by-case basis.

First, assume that $v_i > b_i$. If $x(v_i) > x(b_i)$ or $x(v_i) = x(b_i) = 0$, then clearly there is no incentive to bid b_i. If $x(v_i) < x(b_i)$, we get a contradiction due to Lemma 3. Assume then that $x(v_i) = x(b_i) = 1$. Let $p(v_i) = \frac{R_{-i}}{|\mathcal{W}'|}$ and $p(b_i) = \frac{R_{-i}}{|\mathcal{W}''|}$. From Lemma 5, we know that $\mathcal{W}' = \mathcal{W}''$, therefore $p(v_i) = p(b_i)$ and hence $u_i(v_i) = u_i(b_i)$.

Second, assume that $v_i < b_i$. If $x(v_i) > x(b_i)$ we get a contradiction due to Lemma 3. If $x(v_i) = x(b_i) = 0$, there is no incentive in lying. If $x(v_i) = x(b_i) = 1$, then again, by Lemma 5 we conclude that the payment made by i in both cases are the same, and $u_i(v_i) = u_i(b_i)$. In the case that $x(v_i) < x(b_i)$, let $(R''_i, \mathcal{W}''_i, \mathcal{T}''_i)$ be the critical outcome when i bids b_i. Then, we must have $R''_i \geq R_{-i}$. Since i did not win when bidding v_i, $v_i|\mathcal{W}''| \leq R_{-i}$. But this implies that

$p(b_i) = \frac{R_{-i}}{|\mathcal{W}''|} \geq v_i$, and hence $u_i(b_i) \leq u_i(v_i) = 0$. ■

C. Strategyproofness Over Phases 1 and 2 Combined

We now retrospectively take a look at the two phases of the auction protocol, designed separately, and explain how they work in concert with each other. In the first phase, we compute a revenue-maximizing assignment using Algorithm 3. We then proceed to compute the payments required to ensure this phase is strategyproof using Algorithm 4. Let $\mathbf{q}$ denote the vector of payments as computed by Algorithm 4. The set of agents assigned a channel by Algorithm 3 in the first phase then participate in the second phase, which is the iterative randomized partitioning mechanism shown in Algorithm 2. Let $\mathbf{p}$ be the prices offered to the agents who are winners at the end of Algorithm 2. The auctioneer will allocate channels to this set of agents, and each agent in this set is *charged a final price* of $\max(p_i, q_i)$.

We can show that this two-phase mechanism is truthful in expectation. First, observe that due to the monotonicity of Algorithm 3, no agent that wins in phase one can unilaterally affect the outcome of the set of winning agents in this phase. Combining this with the fact that the payment from phase one is now a minimum payment for agents, we see that there is no incentive at all for an agent to lie in this phase. Furthermore, since phase two is also separately strategyproof, so is the entire mechanism over both phases.

VI. Conclusion

Dynamic spectrum auctions for maximizing revenue were studied in this paper, with interference among secondary users considered. The realistic scenario with no assumption on user valuation distribution is studied. We design a randomized auction for this prior-free setting, and improve the revenue guarantee of previous techniques. Performance analysis of the auction shows that it asymptotically achieves a $\frac{1}{3}$-fraction of the optimal revenue. We then focus on the problem of computing a feasible, interference-free revenue-maximizing channel assignment in a strategyproof fashion. Since the VCG scheme cannot be used in this scenario, we tailor a payment scheme built on Myerson's result on truthful mechanisms for single-parameter agents. The primary user is free to pick any algorithm that solves the maximum K-colourable subgraph problem, yet the strategyproof property of the auction can be maintained.

References

[1] I. Akyildiz, W. Lee, M. Vuran, and S. Mohanty. NeXt generation/dynamic spectrum access/cognitive radio wireless networks: a survey. *Computer Networks*, 50(13):2127–2159, 2006.

[2] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of IEEE Symposium on Foundations of Computer Science (FOCS)*, 2001.

[3] P. Crampton. In *Handbook of Telecommunications Economics, M. Cave, S. Majumdar, I. Vogelsang. (Eds.)*, pages 605–639. Elsevier Science B.V., Amsterdam, 2002.

[4] 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.

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

[6] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

[7] T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.

[8] G. Illing and U. Kluh. *Spectrum auctions and competition in telecommunications*. The MIT Press, 2003.

[9] N. Immorlica, A. Karlin, M. Mahdian, and K. Talwar. Balloon popping with applications to ascending auctions. In *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*, 2007.

[10] Z. Ji and K. Liu. Dynamic Spectrum Sharing: A Game Theoretical Overview. *IEEE Communications Magazine*, 45(5):88–94, 2007.

[11] 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.

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

[13] V. Krishna. *Auction theory*. Academic Press, 2002.

[14] M. Marathe, H. Breu, H. Hunt, S. Ravi, and D. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25(2):59–68, 1995.

[15] M. McHenry and D. McCloskey. New york city spectrum occupancy measurements september 2004. Available from http://www.sharedspectrum.com/inc/content/measurements/nsf/NYC_report.pdf.

[16] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18:511–533, 2001.

[17] R. Myerson. Optimal auction design. *Mathematics of operations research*, pages 58–73, 1981.

[18] N. Nisan and A. Ronen. Algorithmic mechanism design (extended abstract). In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 1999.

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

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

[21] J. Peha. Approaches to spectrum sharing. *IEEE Communications Magazine*, 43(2):10–12, 2005.

[22] S. Sengupta and M. Chatterjee. An economic framework for dynamic spectrum access and service pricing. *IEEE/ACM Transactions on Networking (TON)*, 17(4):1200–1213, 2009.

[23] 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.

[24] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of finance*, pages 8–37, 1961.

[25] X. Zhou, S. Gandhi, S. Suri, and H. Zheng. eBay in the Sky: strategyproof wireless spectrum auctions. In *Proceedings of ACM Mobicom*, 2008.

Appendix

Proof of $\frac{2j\binom{k}{j}\binom{n-k}{n/2-j}}{\binom{n}{n/2}} \geq \frac{j}{2^{k-1}}\binom{k}{j}$ **:**

$$LHS = 2j\binom{k}{j}\frac{(n-k)!}{(n/2-k+j)!(n/2-j)!}\frac{(n/2)!(n/2)!}{n!}$$

$$= 2j\binom{k}{j}(n/2)(n/2-1)\ldots(n/2-k+j+1)$$

$$\frac{(n/2)(n/2-1)\ldots(n/2-j+1)}{n(n-1)\ldots(n-k+1)}$$

$$= 2j\binom{k}{j}(n)(n-2)\ldots(n-2k+2j+2)$$

$$\frac{(n)(n-2)\ldots(n-2j+2)}{2^{k-j}2^j n(n-1)\ldots(n-k+1)} \geq RHS$$

The last $\geq$ holds since k is fixed and odd, and $j \leq \lfloor\frac{k}{2}\rfloor$.