Truthful Spectrum Auctions With Approximate Social-Welfare

Mahmoud Al-Ayyoub and Himanshu Gupta; Stony Brook University, NY

***Abstract*—In cellular networks, a recent trend is to make spectrum access dynamic in the spatial and temporal dimensions, for the sake of efficient utilization of spectrum. In such a model, the spectrum is divided into channels and periodically allocated to competing base stations using an auction-based market mechanism. An efficient auction mechanism is essential to the success of such a dynamic spectrum access model. Two of the key objectives of an efficient auction mechanism are, viz., "truthfulness" (which encourages bidders to truthfully declare their true valuations), and maximizing "social-welfare" (i.e., the total valuation, so that the spectrum is allocated to the bidders who value it the most). Prior works on design of spectrum auction mechanism have only addressed one of the above objectives, and in limited contexts. In this article, we design a spectrum auction mechanism that is truthful *and* yields an allocation that has a social-welfare of within a constant-factor of the optimal. We consider general (pairwise and physical) interference and bidding models. To the best of our knowledge, ours is the first work to design a spectrum auction mechanism satisfying both the above mentioned objectives. We demonstrate the performance of our designed technique through simulations over random and real cellular networks.**

I. Introduction

Usage of wireless spectrum has long been governed by governmental regulatory authorities (e.g., FCC in USA or Ofcom in UK) that divide the spectrum into fixed size chunks to be used strictly for specific purposes, such as broadcast radio/TV, cellular/PCS services, wireless LAN/PANs, etc. This allocation is very long-term and space-time invariant, and is often based on peak usage per provider. Many recent observations have shown that such long-term static allocation of spectrum introduces significant inefficiencies in utilization [1]. To alleviate such inefficiencies, a new policy trend [2] is to make spectrum access dynamic in the spatial and temporal dimensions and responsive to user demands and other considerations. In the case of cellular networks, centralized architectures [1, 3–5] for dynamic spectrum access have gained a lot of interest in the research community due to their practicality and potential impact. In such models, a spectrum broker dynamically allocates spectrum to competing base stations using an auction-based market mechanism. Success of such a model depends on the design of scalable and "efficient" spectrum market mechanisms. Flawed market designs for a precious commodity like spectrum can lead to significant market inefficiencies and adverse economic impacts. This happened in the restructured electricity market in California in 2000 that made international headlines, leading to many academic studies [6–10].

A natural objective of auction-based mechanism is to maximize the generated *revenue* (the sum of the bids or payments by the buyers) [4, 5, 11, 12]. However, such an objective can encourage the spectrum buyers to lie about their real valuations leading to an "untruthful" auction, fear of market manipulation, and indirectly possibly lowered revenue. Moreover, in a competitive environment, buyers may spent a lot of time/effort in predicting the behavior of other buyers and planning against them. Two recent papers address the problem of designing truthful spectrum auctions [13, 14]. Our focus is on auction-based mechanisms that not only encourage truthful behavior but also allocate the spectrum to the bidders who value it the most. The latter goal of maximizing the total valuation is justified in many settings and is extensively studied in economics [15, 16]. Moreover, the bidders with higher valuations are more capable of making good use of the spectrum to build up a viable cellular phone network [17, 18]. Note that, as discussed in Section II-A, it is not possible to design *truthful* auction mechanisms with optimal/approximate revenue.

Problem Addressed. We consider a dynamic auction-based approach to allocate spectrum to competing base stations. The centralized auctioneer acts as the *seller* and the base stations act as the *buyers* of the available spectrum. The items being sold are various channels corresponding to certain (contiguous or non-contiguous) blocks of frequency. The base stations bid for these channels, based on their valuations of these channels.

In the above context, we address the problem of designing a spectrum auction mechanism (i.e., an allocation algorithm) with the following *dual* objectives, viz., (i) encourage truthful behavior from the buyers (i.e., ensure that the buyers "benefit" the most when their bids corresponds to their true valuations), and (ii) at the same time, maximize the "social-welfare," i.e., the total valuation of the allocated channels (by allocating them to the buyers who value them the most).

Closest Prior Works. The above problem of truthful spectrum auction design has been recently addressed in [14] by Zhou et al. under limited interference and bidding models. In particular, they design a spectrum auction mechanism that is truthful, but does not have any performance guarantee on the social-welfare. In other closely related work, Wu et al. [13] focusses on preventing collusion attacks and better revenue (sum of payments from the bidders) in a VCG-like spectrum auction. However, their mechanism requires solving an NP-hard optimization problem, does not guarantee truthfulness, and is limited to simple interference and bidding models.

Our Contributions. In this article, we design a spectrum auction mechanism that yields an allocation (i) that encourages truthful behavior by buyers, and (ii) has an approximate (within a constant-factor of optimal) total valuation. We consider general (pairwise and physical) interference and bidding models. To the best of our knowledge, ours is the first work to design a spectrum auction mechanism satisfying the above dual objectives.

II. Background, Related Work, and Our Contributions

In this section, we present some background material related to our work, and introduce basic terms and definitions from both the spectrum allocation and the auction theory literature. We also discuss related work in more detail, and our contributions.

Dynamic Spectrum Access. In the Dynamic Spectrum Access architectures, the spectrum is allocated dynamically in spatial as well as temporal domain to be more responsive to user demands, and thus, improving utilization. Buddhikot et al. [1], introduced the *coordinated dynamic spectrum access* (CDSA) model for cellular networks. In the CDSA model, there is a centralized entity known as the *spectrum broker* who owns a part of the spectrum called the *coordinated access band* (CAB). The spectrum broker divides the CAB into channels (contiguous or non-contiguous blocks of frequency) and dynamically allocates them to the competing base stations (the buyers) in the region it controls. The base stations express their bids for the available channels using a *bidding function* which specifies the price they are willing to pay for a given set of allocated channels. Periodically, the spectrum broker allocates available channels to the base stations (based on the received bids) under the "wireless interference constraint" such that the total revenue (total price paid by the base stations) is maximized. The above auction-based approach allows the base stations to bid according to the spectrum demands, and the spectrum broker to maximize the revenue generated from allocation of spectrum. However, as mentioned before, this goal is far from ideal in many settings, especially in spectrum auctions, and may cause several problems like market manipulation. To eliminate the fear of market manipulation and allow the bidders to have simple bidding strategies, truthful auction mechanisms are desired.

A. Truthful Auction Mechanisms

In this subsection, we formally define the concepts of auction mechanisms and truthful auction mechanisms. We also discuss VCG auction mechanisms, the only general form of auction mechanism that guarantees truthfulness.

Auction Mechanism. In an auction [16], a set of rational bidders compete over one or more items through a bidding system. An auction is described by the following:

- A finite set O of *allowed outcomes*.
- Each bidder i has a privately-known real function $v_i : O \mapsto \mathbb{R}$ called its *valuation function*, which quantifies the bidder's benefit from each outcome.
- Bidders are asked to declare their valuation functions in the form of *bidding functions* $w = (w_1, \ldots, w_n)$. The bidders may lie about their valuation functions; thus w_i may not be equal to v_i.
- An *auction mechanism* chooses an outcome o based on some criteria over the declared valuation functions.
- In addition to choosing an outcome, the auction mechanism also charges each bidder i a certain amount of money p_i.
- The utility u_i of each bidder i is the difference between its true valuation of the outcome o and its payment p_i, i.e., $u_i = v_i(o) - p_i$. Each bidder's goal is to maximize its utility.

Based on the above model and setting, we define the concepts of auction mechanism and social-welfare.

Definition 1: (Auction Mechanism.) Let O be the set of possible outcomes of an auction. An auction mechanism is a pair of functions (f, p) such that:

- The winner determination function f accepts as input a vector $w = (w_1, \ldots, w_n)$ of bidding (declared valuation) functions and returns an output $f(w) \in O$.
- The payment function $p(w) = (p_1(w), \ldots, p_n(w))$ returns a real vector quantifying the payment charged by the mechanism to each of the bidders.

□

Definition 2: (**Social-Welfare**; Revenue) Social-welfare of an outcome o is defined as the sum of the valuations, i.e., $\sum_i v_i(o)$. Social-welfare may also be defined over declared valuations, i.e., as $\sum_i w_i(o)$.

The revenue of an auction mechanism (f, p) is the sum of the payments $\sum_i p_i(w)$ charged to the bidders for a given declared valuation vector w. □

Generally, the goal of the auction mechanisms is to maximize the total social-welfare, and not necessarily the revenue. The goal, also known as social efficiency, is justified in many settings and is extensively studied in economics [15, 16].

EXAMPLE 1: Let us illustrate the above concepts using the well-known Vickrey's Second-Price Sealed-Bid Auction [19]. Consider an auction wherein a *single* item is up for sale. Each bidder has a certain valuation for the item, and makes a bid accordingly. Here,

- The set of outcomes are $o_1, \ldots, o_n$ where o_i is the outcome in which the item is sold to the i^{th} bidder.
- Valuation function v_i of a bidder i defines the value the bidder assigns to each outcome. Thus, $v_i(o_i)$ is equal to the value of the item for the bidder i, and $v_i(o_j) = 0$ for all $j \neq i$ since a bidder doesn't get any benefit in an outcome where it does not get the item.
- In the Vickery's action mechanism, the item is sold to the bidder with the highest bid, i.e., the winner determination function f chooses an outcome o with maximum social-welfare on w, $\sum_i w_i(o)$. In addition, the payment charged to the highest-bidder is an amount equal to the *second-highest* bid.

Now, consider a two-bidder scenario, where $v_1(o_1) = 5$ and $v_2(o_2) = 20$. Note that $v_1(o_2)$ and $v_2(o_1)$ are zero. Let us assume that the bidders are truthful, i.e., $w_i = v_i$ for all i. Then, the Vickery's mechanism picks the outcome o_2 (i.e., sells the item to the second bidder), and charges the payment of 5 to the second bidder. In this case, the total revenue is 5, and the utilities of the first and second bidders are zero and 15 respectively. □

Truthful Auction Mechanisms. In a selfish environment, bidders may not declare their valuation functions truthfully, if it were to their advantage (result in increase of their utility). Such a behavior may severely damage the resulting welfare and force each bidder to have complex bidding strategies based on its belief/knowledge about the strategies of other bidders. A *truthful* (also known as *incentive-compatible* or *strategy-proof*) mechanism enforces bidders to behave truthfully by offering them incentives (in the form of reduced payments) for such a behavior, or at least, by giving them no incentive for untruthful behavior. These incentives are based on the presumption that each bidder is selfish, and thus, only interested in maximizing its utility. We now formally define the notion of truthful auction mechanism.

Definition 3: (Truthful Auction Mechanisms.) *Given* the valuation functions, in a truthful auction mechanism, each bidder utility is maximized when it truthfully declares its valuation function v_i.

More formally, let the true valuation functions of the bidders be $v = (v_1, \ldots, v_n)$. Consider two declared valuation function vectors, viz., (i) $w = v$ (where each bidder is truthful), and (ii) $w' = (v_1, \ldots, v_{i-1}, w_i, v_{i+1}, \ldots, v_n)$ (where $w_i \neq v_i$; thus all bidders are truthful, except for the bidder i). A mechanism (f, p) is considered *truthful* if $v_i(f(w)) - p_i(w) \geq v_i(f(w')) - p_i(w')$ for all i and $w_i \neq v_i$. □

It is easy to see that Vickery's action mechanism (see Example 1) is indeed truthful. Essentially, by lying about its valuation, a bidder may only hurt its chances of winning, while at the same time, not changing its payment if it wins (since the payment does not depend on its bidding function).

Truthfulness and Revenue Maximization. Note that revenue maximization is not a feasible objective for truthful auction mechanisms. In fact, it is well known that even when auctioning just a single copy of a single item, no truthful mechanism can always attain a guaranteed fraction of the optimal revenue, because there is no way to deal with a single astronomical bidder [18, 20]. In this paper, our main focus is on design of truthful spectrum auction mechanisms, while also maximizing the social-welfare. Through simulations, we show that the revenue of our designed mechanism is close to that delivered by the best-known approximation algorithm, and is an order of magnitude better than a naive truthful spectrum auction.

VCG Auction Mechanisms. The only general mechanism[1] that guarantees truthfulness is due to Vickrey-Clarke-Groves (VCG) [19, 21, 22], and in some scenarios, it is known that no other method exists [16]. In restricted settings, however, other approaches [23, 24] may exist. Informally, the celebrated VCG mechanism finds the outcome o with maximum social-welfare, and charges each winner i an amount equal to the total "damage" that it causes to the other bidders, i.e., the difference between the social-welfare of the whole system with and without i's participation [15]. The key property of such a mechanism is that it "aligns" each bidder's utility with the social-welfare [15]. Thus, by maximizing its own utility, each bidder is also maximizing the social-welfare. Below, we give a formal definition of the VCG mechanism.

Definition 4: (VCG Mechanism.) A VCG mechanism is an auction mechanism (f, p) (see Definition 1) that satisfies the following two conditions, for any given declared valuation functions $w = (w_1, \ldots, w_n)$.

- $f(w) \in \arg\max_o \sum_i w_i(o)$, i.e., the winner determination function f chooses an outcome that maximizes the social-welfare according to w.
- The payment functions are determined by the VCG formula $p_i(w) = (-\sum_{j \neq i} w_i(f(w))) + h_i(w_{-i})$, where each $h_i(w_{-i})$ is an arbitrary function of $w_{-i} = (w_1, \ldots, w_{i-1}, w_{i+1}, \ldots, w_n)$. For non-negative declared valuations, the function $h_i(.)$ is usually chosen according to the Clarke pivot rule which suggests the choice of $h_i(w_{-i})$ as the maximum social-welfare due to w_{-i}, i.e., $h_i(w_{-i}) = \max_{o \in O} \sum_{j \neq i} w_i(o)$.[2]

□

One of the main shortcomings of VCG mechanisms is that they may result in low (even zero) revenue in some cases. But, VCG's payment function is key to ensuring its truthfulness, and altering its payment scheme may destroy its truthfulness property.

VCG Mechanisms for NP-hard problems. Unfortunately, a VCG mechanism requires solving an optimization problem of maximizing social-welfare, which can be NP-hard in many settings. Furthermore, Nisan and Ronen [16] showed that choosing an allocation with even approximate social-welfare destroys the truthfulness of the mechanism. However, to circumvent the above obstacle, they introduce maximal-in-range (MIR) mechanisms (formally defined in the next subsection) where a suboptimal allocation can be used while maintaining the truthfulness property.

B. Related Works, Our Approach and Contributions

In this section, we discuss related works, our approach to design of truthful spectrum auction mechanisms, and the main contributions of our work.

Related Works. Traditional auction mechanism are not directly applicable to spectrum auctions due to the "multi-winner" property of each item (due to spatial reuse of spectrum channels) and wireless interference constraints. Moreover, the corresponding optimization problem of maximizing social-welfare in the context of spectrum auctions is known to be NP-hard [11], which makes VCG auction mechanisms inapplicable. Since the truthfulness property is key to our work, below we discuss recent works on truthful spectrum auctions in detail.

Truthful Spectrum Auctions. To the best of our knowledge, there has been only one work till date, viz., [14], that has

[1] Vickrey's Second-Price Sealed-Bid Auction is only for simple single-item auctions, and looses its truthfulness when generalized to more complicated multi-item auctions.

[2] Intuitively, to maximize its revenue, the mechanism should choose the maximum payment a winning bidder can pay without "eating up" all its utility gain [15]. This is achieved by the Clarke pivot rule.

designed truthful mechanisms for spectrum auction. The truthful mechanism designed by Zhou et al. [14] however does not address the goal of maximizing social-welfare. Moreover, their approach is limited to only simple (single-minded or range) bidding functions and pairwise interference model. As observed in [14, 18], it is rather straight-forward to design a truthful auction mechanism without any regard for social-welfare. But, the authors in [14] show through simulations that their mechanism returns better social-welfare and revenue compared to a simple truthful mechanism. Recently, this work has been extended to consider double auctions [25].

In another closely related work, Wu et al. [13] design a spectrum auction mechanism based on VCG mechanism. They focus on modifying the VCG payment function to eliminate colluding attacks by losing bidders and to improve the total revenue. However, their altered payment scheme destroys the truthfulness property of the VCG scheme. In addition, their mechanism requires solving an integer linear programming (NP-hard) problem, which makes their approach impractical for large networks. Note that in practice cellular networks may have thousands of nodes [26]. Finally, they assume either a single-channel system or that each bidder is interested in only one channel in a multi-channel system.

Other Spectrum Auction Works. Recently there has been lots of works on dynamic spectrum allocation using auctions, but most of the works have focussed on the goal of revenue maximization [4, 5, 11, 12] without worrying about the truthfulness of the auction mechanism. However, as mentioned before, an untruthful auction mechanism can encourage the bidders to lie about their valuations, which may lead to market manipulation and lowered revenues. In addition, in a competitive environment, buyers may be forced to spent a lot of time/effort in predicting the behavior of other buyers and planning against them.

Our Approach: Truthfulness with Approximate Social-Welfare. In this article, we focus on designing a spectrum auction mechanism that is truthful and selects an outcome with approximate social-welfare, for general interference models and bidding functions. In our approach, we make use of the recent result by Dobzinski and Nisan [27] on design of truthful auction mechanisms with approximate social-welfare for multi-unit auctions (MUA) using a maximal-in-range (MIR) mechanism. Thus, we start with formally defining MIR mechanisms and multi-unit auctions. Then, we give an outline of our approach.

Maximal-in-Range Mechanisms. In [16], the authors provide a computationally-efficient way to overcome the problem of finding an allocation with *optimal* social-welfare in VCG mechanisms, since finding such an allocation may be NP-hard in some cases. In particular, they show that an auction mechanism is truthful if it (i) chooses an outcome that optimizes social-welfare over a fixed *subset* of the outcomes, and (ii) uses VCG payments (as defined in Definition 4). Such mechanisms are termed *Maximal-In-Range (MIR)* and formally defined below.

Definition 5: (Maximal-In-Range (MIR) Mechanism.) Let V_i be the set of all possible valuation functions of bidder i, and $V = \prod_i^n V_i$ be the space of all possible valuation functions. Let $\mathcal{O}$ denote the range of the winner determination function f at V, i.e., $\mathcal{O} = \{f(v) | v \in V\}$. We say that f is *maximal in its range* if for every $v \in V$, $f(v)$ maximizes the social-welfare over $\mathcal{O}$. □

Multi-Unit Auctions (MUA). Multi-unit auctions (MUA) have been heavily studied in economics due to its practical implications. In a MUA, a set of m identical items are up for auction among bidders, and each bidder expresses interest for certain *quantities* of the items, without any preference to any specific item. Thus, the valuation function of a bidder i can be represented[3] as $v_i : \{1, \ldots, m\} \mapsto \mathbb{R}$, where $v_i(q)$ is the value for obtaining q items. In [27], the authors design an MIR mechanism for multi-unit auctions that is truthful and yields an allocation with approximate social-welfare.

Our Approach. As mentioned above, traditional auction mechanisms cannot be directly used on spectrum auctions due to the "multi-winner" property of each auctioned item in the spectrum auctions. Our approach utilizes the geographical nature of the spectrum auction problem to re-formulate it as a *set* of multi-unit auction instances. Then, we use the MIR mechanisms for multi-unit actions from [27] to solve each instance, i.e., independently determine spectrum allocation with approximate social-welfare for each instance. We ensure truthfulness by using VCG payments. Finally, we combine the allocations over these independent instances in a way that preserves truthfulness and the approximation ratios of the social-welfare.

Our Contributions. Basically, we present a simple and general approach to dynamically allocate spectrum to competing base stations with the goal of maximizing the social-welfare while maintaining truthfulness. We consider general bidding functions and interference models. For the pairwise interference model, we consider the unit-disk, non-uniform disk, and the pseudo-disk models. For the physical interference model, we consider uniform as well as non-uniform power transmission models. Our contributions can be summarized as follows:

- For the pairwise interference with unit-disk model and general bidding functions, we present a truthful auction mechanism that yields an allocation whose social-welfare is within a constant-factor of the optimal. We extend the truthful mechanism and its approximation result to (i) non-uniform disk and pseudo-disk pairwise interference models, (ii) k-minded bidding function (where the bidder expresses its valuations for at most k quantities of channels), (iii) non-orthogonal channels, and (iv) multi-type channel auctions.
- For the physical interference model with uniform power transmissions and general bidding functions, we present a truthful auction mechanism that yields an allocation whose social-welfare is within a constant-factor of the optimal. The result is extended to k-minded bidding functions, and non-uniform power transmission model.

[3]Note that such a representation can be easily mapped to the original form wherein a valuation function maps outcomes to real numbers.

Revenue Maximization. We note that our developed techniques can also be used to design an approximation algorithm for the revenue maximization problem [11] in spectrum auctions. In particular, by charging each bidder a payment equal to their declared valuation for the allocated number of channels, our algorithm's allocation has a revenue of within a constant-factor of the optimal revenue possible.[4] This improves the prior best-known result by Subramanian et al. [11] for single-type channels, by providing an approximation algorithm for even more general (non-complementary) bidding functions.

III. Spectrum Auction Under Pairwise Interference

In this section, we address our problem of designing truthful spectrum auctions with approximate social-welfare under the pairwise interference model. In the first subsection, we will consider the unit-disk model, wherein the coverage region of each base station is assumed to be a disk of uniform radius. In later subsections, we extend our techniques to non-uniform disk and pseudo-disk interference models. We start with defining our network model, the core concepts, and formulating the problem formally.

Network Model; Interference Graph. Our model of a cellular network consists of a set of geographically distributed base stations. Spectrum is divided into *orthogonal* channels of the *same type*, and the spectrum auction involves each base station bidding for certain *quantities* of channels (as in basic multi-unit auctions). In a later subsection, we will consider non-orthogonal channels and multi-type channel networks.

Each base station is associated with a region around it called its *cell*; each base stations serves its clients in its cell. To communicate, the base station and the client must operate "interference-free" on the same channel. In cellular networks, wireless interference at a client may arise due to multiple near-by base stations operating on the same channel. In the pairwise interference model, pairs of base stations with intersecting cells are said to *interfere* with each other if operating on the same channels. These pairs of interfering base stations can be represented by simple edges over base stations as vertices in an interference graph, as defined below.

Definition 6: (Interference Graph G_t.) The interference graph $G_t = (N_t, E_t)$ is an undirected graph where each vertex represents a base station and there is an edge $(i, j) \in E_t$ between i and j if the corresponding base stations "interfere".

As mentioned before, two base stations are said to interfere when their corresponding cells intersect. Note that interfering base stations should not be allocated a common channel. □

Valid Spectrum Allocation. Informally, our spectrum auction problem is to allocate channels to base stations so as to maximize the social-welfare and maintain truthfulness. However, the allocation of channels should be done without violating the interference constraints. We formalize this by defining a concept of valid spectrum allocation [11], which essentially represent the possible outcomes of the spectrum auction.

[4] In the revenue maximization problem, we assume that a bidder is not charged more than its bid for the allocated number of channels.

Definition 7: (Valid Spectrum Allocation.) A *spectrum allocation* is a set of (base station, channel) pairs, i.e., a spectrum allocation is a set

$$\{(i, c) | i \text{ is a base station}, c \text{ is a channel}\}.$$

A spectrum allocation $\mathcal{A}$ is considered *valid* if no two interfering base stations are associated with a common channel in $\mathcal{A}$. More formally, $\mathcal{A}$ is valid, if for all (i, c) and (j, c) in $\mathcal{A}$, i and j do not interfere. □

Representation of Valuation and Bidding Functions. In a spectrum auction of channels of same type, a bidder i's valuation of an outcome/allocation o depends only on the number of channels i is getting in o. Thus, we represent bidder i's valuation function v_i as $v_i : \{1, \ldots, m\} \mapsto \mathbb{R}$, where m is the total number of channels and $v_i(q)$ denotes bidder i's value for obtaining q channels. Recall that the bidding function w_i for a bidders i is a declaration of its privately-known valuation function v_i. Thus, the bidding function w_i is represented similarly as $w_i : \{1, \ldots, m\} \mapsto \mathbb{R}$. We assume free disposal (i.e., valuation for higher number of channels is larger than smaller number of channels), and that valuation of zero channels is zero.

General-Minded and k-minded Bidding Functions. In the most general model, a bidder has a valuation for any number of channels, and thus, the bidding functions is represented by m real numbers – one for each quantity of channels. For efficiency and practicality issues, another model is commonly assumed in the literature, viz., the k-minded bidding function, wherein the bidder expresses its valuations for at most k quantities of channels.

TSA-MSW (Truthful Spectrum Auctions with Maximum Social-Welfare) Problem. Given an interference graph, *number* of channels, and the bidding functions for the base stations, the *TSA-MSW problem* is to design a truthful auction mechanism that returns a valid spectrum allocation with maximum social-welfare.

Thus, the TSA-MSW problem involves determining (i) a *valid* spectrum allocation with optimal social-welfare, and (ii) payments by each bidder, so that the overall mechanism is truthful. TSA-MSW problem is NP-hard even without the truthfulness objective [11]. Thus, we focus on designing a mechanism that is truthful and yields a valid spectrum allocation with approximate social-welfare.

Input and Output Sizes. If n and m denote the number of nodes and channels respectively, then note that the size of the input is polynomial in n and $\log m$ (since the input only includes the *number* of channels). On the other hand, the size of the output as defined in Definition 7 may be polynomial in n and m, and thus, exponential in the input size. However, it is easy to modify Definition 7 so that valid spectrum allocation is polynomial in n and $\log m$, by assocaiting a *number* of channels with each node. We have defined valid spectrum allocation as in Definition 7 for simplicity of presentation.

A. TSA-MSW Problem in Unit-Disk Model

In the *unit-disk* model, the coverage region of each base station is assumed to be a disk of uniform radius d. For simplicity

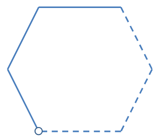

Fig. 1. To ensure validity of our algorithms, the hexagons we use are open from side and closed from the other. For simplicity, we will use closed hexagons in the remaining figures.

of presentation, we assume distances to be normalized, i.e., $d = 1$. Thus, two base stations interfere if they are within a unit distance from each other. We start with giving an outline of our allocation algorithm (i.e., the winner determination function) for the unit-disk model. Then, we will discuss various parts of the algorithm in detail, and prove its truthfulness and social-welfare approximation ratio.

Outline of the Allocation Algorithm. Our approach utilizes the geographical nature of the spectrum auction problem to divide it into smaller and more tractable subproblems. Then, we solve each subproblem independently and "combine" the allocations, without sacrificing much on the overall approximation factor of the final social-welfare. At a high-level, our algorithm consists of the follows steps.

1) Divide the entire network region into small hexagons[5] of side-length one unit each. See Figure 1. This division ensures that any pair of base stations in the same hexagon interfere with each other (due to the unit-disk interference model).
2) Uniformly-color the hexagons with enough colors, such that base stations in co-colored hexagons are more than two-unit distance away and hence do not interfere.
3) Allocate channels to base stations in each hexagon *independently*, treating it as a multi-unit auction (MUA) and using techniques similar to [27]. Note that the interference subgraph in each hexagon is actually a complete graph.
4) For each color, combine the results from all hexagons of that color.
5) Pick the color that has the highest total social-welfare.

The above gives a (7γ)-approximate solution, where γ is the approximation factor in Step (3), and 7 is the number of colors used to color the hexagons. The value of γ is 2 for general-minded bidding, and $(1 + \epsilon)$ for k-minded bidding functions for any $\epsilon > 0$. Furthermore, the above algorithm, and all the generalizations discussed afterwards, run in time polynomial in the size of the input, i.e., in n, the number of nodes, and $\log m$, where m is the number of channels.

Hexagonal Division, and 7-Coloring of Hexagons. Our Algorithm starts by diving the plane into hexagons of side-length one unit each (creating a hexagonal division of the plane), and proceed to uniformly coloring these hexagons using 7 colors. See Figure 2. In such a coloring, the following two properties hold.

(P1) Every pair of base stations in the same hexagon interfere with each other (i.e., are connected by an edge in the interference graph).

[5]These hexagonal division is not to be confused with the actual *cells* associated with the base stations.

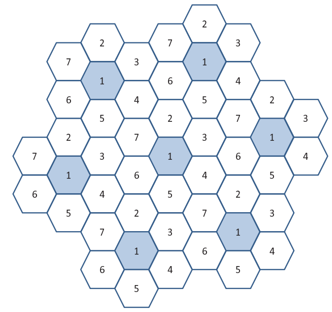

Fig. 2. Hexagons uniformly-colored using 7 colors.

(P2) Base stations in different hexagons with same color do not interfere with each other (i.e., are *not* connected by an edge in the interference graph).

Property (P1) follows directly from the definition of unit-disk interference, while Property (P2) follows from the fact that the distance between base stations in different hexagons with the same color will be at least $(\sqrt{3(7)} - 2) > 2$.

Allocation in Each Hexagon. The above properties imply that the channels cannot be re-used inside the same hexagon, but can be fully re-used across different hexagons of the same color. Thus, allocation in each hexagon can be treated as an MUA, and using techniques similar to [27] we can design an MIR mechanism with approximate social-welfare. Below, we describe their technique in detail for general-minded and k-minded functions.

General-Minded Bidding Model. In case of general-minded bidding model, the available m channels are split into N_H^2 bundles of size $\left\lfloor \frac{m}{N_H^2} \right\rfloor$ each, where N_H is the number of base stations in hexagon H, and a single bundle of remaining channels. Using dynamic programming, we can *optimally* allocate these bundles to the N_H bidders in time polynomial in N_H. The above approach yields an allocation whose social-welfare is at least 1/2 of the optimal possible (see [27]).

k-minded Bidding Functions. In the case of k-minded bidding function, a restricted form of allocation known as the t-*round* allocation is used. For a given t, a t-round allocation allocates l $(l \leq m)$ channels to a subset T of the bidders where $|T| \leq t$; this part of the allocation is done optimally by exhaustive search for each l and T. Also, for each l and T, the remaining $(m - l)$ channels are divided into equi-sized bundles and distributed optimally to the remaining bidders, as in the case of general-minded bidding model. Finally, the best allocation among the mt such allocations is picked as the optimal t-round allocation. The above allocation algorithm runs in polynomial-time (in number of base stations and k) for a fixed t, and yields a $\left(1 - \frac{1}{t+1}\right)$-approximate allocation [27].

Combining The Results. Since base stations in different hexagons of same color do not interfere with each other (Property (P2)), we can combine allocations of co-colored hexagons to form one single allocation. Thus, we get seven allocations, one for each color. Among these seven allocations, we pick the allocation with the highest social-welfare, as our final solution.

Proof of Truthfulness and Approximability. Our overall spectrum auction mechanism consists of the above described allocation algorithm (winner determination function) combined with VCG payments (as described in Definition 4). In the below theorem, we prove that this overall auction mechanism is truthful and returns a valid spectrum allocation with approximate social-welfare.

Theorem 1: For the TSA-MSW problem under the pairwise interference with unit-disk model, the above described auction mechanism is truthful and returns a valid spectrum allocation whose social-welfare is 14-approximate for the general-minded bidding model and is $7(1+\epsilon)$-approximate for the k-minded bidding model for a given $\epsilon > 0$.

Proof: Truthfulness. Our allocation algorithm picks a t-round allocation with the highest social-welfare, for a given t (for the case of general-bidding functions, t can be considered to be zero). Thus, our allocation algorithm is maximal in its range, where the range of allocations/outcomes is restricted to t-round allocations. Thus, our auction mechanism is truthful since MIR allocations with VCG payments are truthful [16].

Approximate Social-Welfare. First, note that by the properties (P1) and (P2) of the hexagonal division, the allocation returned by our algorithm is valid. Now, let us prove the approximation factor for the general-minded bidding model; the proof for k-minded bidding model is similar. Consider a particular color c, and for the set of all hexagons colored c, let I_c be the allocation constructed by our algorithm and O_c be the allocation with optimal social-welfare. We show that the social-welfare of I_c is within a factor of 2 of that of O_c. Note that, for any particular hexagon cell, our algorithm constructs an allocation whose social-welfare is within a factor of 2 of the optimal for that hexagon. Since I_c's (O_c's) social-welfare is the sum of the social-welfares of the constructed (optimal) allocations for the individual c-colored hexagons, we get that the social-welfare of I_c is within a factor of 2 of that of O_c. Now, since there are seven colors and we pick the best of the seven allocations, the social-welfare of the returned allocation is within a factor of 14 of the overall optimal social-welfare. ∎

B. TSA-MSW Problem in Non-Uniform Disks Model

We now extend our techniques of previous subsection to the *non-uniform disk* model, wherein cells of the base stations are disks of possibly different radii. As before, two base stations are considered to interfere if their cells intersect. Let the maximum and the minimum disk radii in the network be $d_{\max}$ and $d_{\min}$ respectively. For simplicity of presentation, we assume that the distances are normalized, i.e., $d_{\min} = 1$.

For the above disk model, we divide the base stations into classes depending upon their cell's radius, and then solve the spectrum allocation problem for each class independently. Finally, we pick the allocation of the class that has the highest social-welfare. Thus, our algorithm consists of the following steps.

1) Classify the base stations into $\lfloor \log(d_{\max}) \rfloor$ *radius-classes*, based on their cell's radius. In particular, class L contains base stations whose cell's radius d_i lie in the range $d_i \in [2^L, 2^{L+1})$.
2) For each radius-class L:

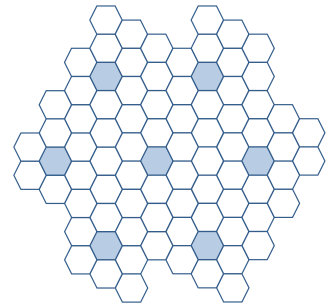

Fig. 3. Hexagons uniformly-colored using 12 colors.

   a) Divide the network region into hexagons of side-length 2^L each.
   b) Uniformly-color the hexagons using 12 colors as shown in Figure 3.
   c) Independently, for each hexagon H, allocate channels to base stations of radius-class L contained in H. The interference subgraph induced by these base stations is a complete graph, and thus, we can use the same technique as for the unit-disk model.
   d) For each color, combine the results from all hexagons of that color. Note that base stations of class L in different co-colored hexagons do not interfere with each other.
   e) Pick the color that has the highest total social-welfare.
3) The above gives an allocation for each radius-class. Pick the allocation for the radius-class that has the highest social-welfare.

Theorem 2: For the TSA-MSW problem under non-uniform disks interference model, the auction mechanism based on the above allocation algorithm and VCG payments is truthful and returns a valid spectrum allocation whose social-welfare is $24\lfloor \log(d_{\max}) \rfloor$-approximate for the general-minded bidding model and is $12\lfloor \log(d_{\max}) \rfloor (1+\epsilon)$-approximate for the k-minded bidding model for a given $\epsilon > 0$.

Proof: The truthfulness of the auction mechanism follows from the same arguments as in the proof of Theorem 1.

Validity. Validity of the returned allocation follows from the fact that for each radius-class L, the base stations of the radius class L satisfy the Property (P2) of previous subsection. Note that with 12-coloring of hexagons of side-length 2^L each, the distance between base stations in different co-colored hexagons is at least $(\sqrt{3(12)}-2)2^L = 2(2^{L+1})$, which is not close enough to create interference between base stations in radius-class L.

Approximate Social-Welfare. The proof of approximation follows from similar arguments as in the proof of Theorem 1, except for the fact that we use 12 colors here (instead of 7) and the extra $\lfloor \log(d_{\max}) \rfloor$ factor comes due to the number of radius-classes considered independently. ∎

C. TSA-MSW Problem in Pseudo-Disk Model

We now extend our techniques to the most general case of *pseudo-disk model* model, wherein cells of the base stations have irregular shapes but are contained within a disk of radius

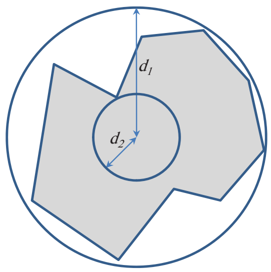

Fig. 4. An example of the pseudo-disk model.

d_1 while containing a disk of radius $d_2 \leq d_1$. See Figure 4. For simplicity of presentation, we assume that d_1 and d_2 are the same for all base stations. Techniques of previous subsection can be used to extend our results below to the case wherein d_1 and/or d_2 may be different for different cells. Also, for clarity of presentation, we use $d_2 = 1$.

The allocation algorithm for the pseudo-disk model is similar to the one for unit-disk model, except that the side-length of the hexagons and the coloring scheme are different. To ensure the correctness of the unit-disk approach in the context of pseudo-disk model, we need to do the division and coloring appropriately to ensure that Properties (P1) and (P2) of Subsection III-A hold. To ensure Property (P1), we divide the network region into hexagons of side-length one unit, as in the case of unit-disk model. Below, we compute the number of colors required to *uniformly* color the hexagons, in order to satisfy Property (P2).

Required Number of Colors. To satisfy Property (P2), i.e., to ensure that base stations in different hexagons with the same color do not interfere, we must color the hexagons in a way that the distance between any two points in different hexagons of the same color is greater than $2d_1$. To estimate the number of colors required, we make use of the following two lemmas from [26, 28].

Lemma 1: In a hexagonal division with side-length s and uniformly-colored[6] with x colors, the distance between the centers of two hexagons of the same color is at least $\sqrt{3x}s$. □

Lemma 2: A hexagonal division can be uniformly colored using c colors if and only if c is of the form $i^2 + j^2 + ij$ for some positive integers i and j. □

Now, by Lemma 1 above, to ensure a distance of $2d_1$ between co-colored hexagons, the number of colors must be greater than $4d_1^2/3$. Then, by Lemma 2 above, the minimum number of colors required would be given by:

$$q_w = \min\{x | x \geq 4d_1^2/3 \text{ and } x = i^2+j^2+ij \text{ where } i,j \in \mathbb{Z}^+\}.$$

Approximate Social-Welfare. Using arguments similar to before, the above hexagonal division and coloring yields an allocation algorithm with approximation factor of $2q_w$ and $q_w(1+\epsilon)$ for the general-minded and k-minded bidding models respectively.

D. Non-Orthogonal Channels; Multiple Types of Channels

Thus far, we have assumed that the channels in our network model are orthogonal and of the same type. Here, we discuss relaxation of these two assumptions.

Non-Orthogonal Channels. In general, the non-orthogonal (overlapping) nature of channels can be modeled using a channel graph G_c over channels as vertices, wherein there is an edge between two channels c_i and c_j if they are non-orthogonal. Let I be a maximum independent set in G_c. If we can somehow compute I, then we can just use I as the set of channel to allocate in *each* hexagon of our technique; this will maintain our approximation ratios because of the following two facts. First, reuse of I in different hexagons can be done without any changes to our techniques, since either they have no interference edges between them or they are colored differently. Second, within any hexagon, using I is sufficient, since all the channels are of the same type and no two base stations within a hexagon can share a channel. Thus, computation of a maximum independent set in G_c is sufficient to use of our techniques and preserve the approximation guarantees. Now, if the channels are contiguous blocks of spectrum, then the channel graph G_c is an "interval graph" wherein the maximum independent set (MIS) can be easily computed using a simple greedy approach. For non-contiguous channels, our techniques can still be used, but the performance guarantees do not hold.

Orthogonal Channels of Multiple Types. We have assumed till now that the available channels are of the same type; thus, each base station had valuations for *quantities* of channels without any preference for specific channels. Such a model allowed us to treat the allocation subproblems in each hexagon as multi-unit auctions (MUA). However, if we have multiple types of channels, and bidders have valuations for various *sets* of channels, then the allocation subproblem in each hexagon must be instead treated as a Combinatorial Auction (CA) problem, which is more difficult than the MUA problem. For the CA problem, the best known result is by Holzman et al. [29] who give a truthful MIR mechanism that achieves $O(\frac{m}{\sqrt{\log m}})$ of the maximum social-welfare for general-minded bidding model, where m is the total number of items. Thus, for spectrum auction of multi-type channels, we can use [29]'s allocation algorithm to allocate channels in each hexagon, and combine results as before. The resulting auction mechanism will be truthful and yield a $O(\frac{m}{\sqrt{\log m}})$-approximate social-welfare, where m is the total number of channels.

IV. TSA-MSW Problem Under Physical Interference

In this section, we extend our techniques from the previous section to the case of physical interference. We start with considering the uniform-power transmission model.

A. Uniform Transmission Powers

In this subsection, we assume that each base station operates using the same transmission power P; we relax this assumption in the next subsection. We start by introducing the physical interference model, and redefining the concept of valid spectrum allocation in this context.

Physical Interference Model. In the physical interference model, a reception at a certain distance from a base station

[6]Informally, in a uniform-coloring of hexagons, the distance between the "closest" hexagons with the same color is uniform.

is successful, if the "signal to noise plus interference ratio" (SINR) at the receiver is greater than a threshold β. More formally, a reception from a base station i is successful at a point p if and only if,

$$\frac{P/d_i^\alpha}{\mathcal{N} + \sum_{j \in B'} P/d_j^\alpha} \geq \beta \tag{1}$$

where B' is the set of other base stations operating on the same channel as i, d_x is the distance of the point p from a base station x, $\mathcal{N}$ is the background noise, and α is the path loss exponent based on the terrain propagation model.

Communication Radius (r). The communication radius [11] r of a base station i is the maximum distance from i within which we *want* the SINR from i to be at least as large as β. Essentially, the above is based on the stipulation that the cell of base station i is a disk of radius r. In our context, the value of r can be arbitrarily large (but finite), since the approximation ratio and time complexity of our designed algorithms are independent of r. Thus, the concept of communication radius must not be looked upon as an assumption. We assume uniform communication radius; non-uniform communication radii can be handled in the similar manner as non-uniform disks in Section III-B.

Valid Spectrum Allocation. In the context of physical interference model, a spectrum allocation $\mathcal{A}$ is considered valid if it satisfies the following condition. For each pair (i, c) in $\mathcal{A}$, let $B_{i,c}$ denote the set of base stations that have been allocated the channel c in $\mathcal{A}$. Now, for $\mathcal{A}$ to be valid, for every (i, c) in $\mathcal{A}$ and every point p within a distance of r from i, SINR at p due to i and c should be greater than β; i.e, the following should hold:

$$\frac{P/d_i^\alpha}{\mathcal{N} + \sum_{j \in B_{i,c}} P/d_j^\alpha} \geq \beta$$

where d_x is the distance of base station x from the point p.

Allocation Algorithm. The allocation algorithm for physical interference model is similar to the one for pairwise interference model in the previous section, except for the chosen side-length of the hexagons and the number of colors used for uniform-coloring of the hexagons. To ensure correctness of our approach in the context of physical interference, we need to do the hexagonal division and coloring in such a way that the following two properties are satisfied.

(P'1) Every pair of base stations b_1 and b_2 in the same hexagon must "interfere" with each other when operating on the same channel, even if no other base station is active. In other words, no valid spectrum allocation must assign the same channel to b_1 and b_2.

(P'2) If in each hexagon with the same color there is at most one active base station, then the transmission from each of these base stations must be successful within their communication radius.

To ensure Property (P'1), we can just divide the network region into hexagons of side-length

$$R = \frac{(\sqrt[\alpha]{\beta} + 1)r}{2}.$$

It is easy to see that Property (P'1) is satisfied for the above hexagonal division.

Coloring Hexagons to Satisfy Property (P'2). In the below Lemma 3, we will show that Property (P'2) can be satisfied by ensuring that the minimum distance between hexagons with the same color is at least $\sqrt{3q_1'}R$, where R is as defined above and q_1' is as defined below.

$$q_1' = \left(\frac{4\sqrt{7}}{(3\sqrt{7}-6)(\sqrt[\alpha]{\beta}+1)}\right)^2 \left(\frac{6\beta}{(\alpha-2)}\right)^{\frac{2}{\alpha}}.$$

Then, using arguments similar to Section III-C, the number of colors required to satisfy Property (P'2) is given by:

$$q_1 = \min\{x | x \geq q_1', x \geq 7, \text{ and } x = i^2 + j^2 + ij \text{ where } i, j \in \mathbb{Z}^+\} \tag{2}$$

Note that in our context we should use at least 7 colors, irrespective of the values of α and β. We now state and prove the Lemma 3 used in the above argument.

Lemma 3: Property (P'2) is satisfied if the minimum distance between hexagons with the same color is at least $\sqrt{3q_1'}R$, where R and q_1' are as defined above.

Proof: Consider a base station i in a hexagon H of color c. Partition all c-colored hexagons surrounding H into hierarchical levels. In a uniform-coloring, the first level will contain 6 hexagons of color c and each such hexagon H' is at distance[7] of at least $(\sqrt{3q_1} - 2)R$ from H (from Lemma 1). Similarly, the second level contains 12 hexagons at a distance of at least $(3\sqrt{q_1} - 2)R$ from H. In general, the l^{th} level contains $6l$ hexagons at a distance of at least $(\frac{3}{2}\sqrt{q_1}l - 2)R$ from H.

Now consider a point p within the communication radius r from the base station i. Then, the total signal received at the point p due to all other base stations (at most one per c-colored hexagon) active on the same channel as i is at most:

$$\begin{aligned}
& \sum_{l=1}^{\infty} 6l \cdot \frac{P}{\left(\frac{(\frac{3}{2}\sqrt{q_1}l-2)(\sqrt[\alpha]{\beta}+1)}{2} - 1\right)^\alpha r^\alpha} \\
\leq \quad & \frac{6P}{r^\alpha} \sum_{l=1}^{\infty} \frac{l}{\left(\frac{l\sqrt{q_1}(3\sqrt{7}-4)(\sqrt[\alpha]{\beta}+1)}{4\sqrt{7}} - 1\right)^\alpha} \\
\leq \quad & \frac{6P}{r^\alpha} \sum_{l=1}^{\infty} \frac{l}{\left(\frac{l\sqrt{q_1}(3\sqrt{7}-6)(\sqrt[\alpha]{\beta}+1)}{4\sqrt{7}}\right)^\alpha} \\
= \quad & \frac{6P}{\alpha-2} \cdot \left(\frac{4\sqrt{7}}{\sqrt{q_1}(3\sqrt{7}-6)(\sqrt[\alpha]{\beta}+1)r}\right)^\alpha.
\end{aligned}$$

Above, the second equation follows from the following two facts: (i) $\frac{3}{2}\sqrt{q_1}l \geq 3\sqrt{7}/2$ (since $q_1 \geq 7$), and (ii) for $x \geq 3\sqrt{7}/2$, we have $(x-2) \geq \frac{x}{3\sqrt{7}/(3\sqrt{7}-4)}$. And, the third equation follows from the following facts: (i) $l\sqrt{q_1}(3\sqrt{7} - 4)(\sqrt[\alpha]{\beta}+1)/4\sqrt{7} \geq (3\sqrt{7}-4)/2$, since $q_1 \geq 7$ and $\sqrt[\alpha]{\beta} \geq 1$, and (ii) for $x \geq (3\sqrt{7}-4)/2$, $(x-1) \geq \frac{x}{(3\sqrt{7}-4)/(3\sqrt{7}-6)}$.

For simplicity, we assume ambient noise to be zero; non-zero noise can be incorporated using techniques similar

[7] By distance between two hexagons we mean that the distance between *any* point in H' and *any* point in H.

to [30]. Now, using the value of q_1 from Equation 2, the SINR at point p due to the transmitted at base station i is at least:

$$\frac{P}{r^\alpha} \cdot \frac{\alpha - 2}{6P} \cdot \left(\frac{\sqrt{q_1}(3\sqrt{7} - 6)(\sqrt[\alpha]{\beta} + 1)r}{4\sqrt{7}} \right)^\alpha \geq \beta$$

∎

Overall Allocation Algorithm. In the above paragraph, we discussed hexagonal division and its coloring in a uniform way so as to satisfy Properties (P'1) and (P'2). Now, note that Property (P'1) ensures that allocation in each hexagon can be treated as a multi-unit auction, while Property (P'2) allows us to re-use channels across different hexagons with same color. Thus, we can use the same allocation algorithm as for the unit-disk model, with the above hexagonal division and coloring. Use of VCG payments yields an overall truthful auction mechanism. Thus, we have the following theorem.

Theorem 3: For the TSA-MSW problem under the physical interference model with uniform transmission power, the auction mechanism based on the above described allocation algorithm and VCG payments is truthful and returns a valid spectrum allocation whose social-welfare is $2q_1$-approximate for the general-minded bidding model and is $q_1(1+\epsilon)$-approximate for the k-minded bidding model for a given $\epsilon > 0$. Here, q_1 is as defined in Equation 2. □

B. Non-Uniform Transmission Power

We now consider the model wherein different base stations may operate on different transmission powers. Let the maximum and the minimum transmission power levels used in the network be $P_{\max}$ and $P_{\min}$ respectively. For simplicity of presentation, we assume that transmission powers are normalized, i.e., $P_{\min} = 1$. For this non-uniform transmission model, the physical interference model and valid spectrum allocations can be appropriately defined. For simplicity, we assume uniform communication radius; *non-uniform communication radii can be handled in the similar manner as non-uniform disks in Section III-B.*

Allocation by Division into Power Classes. Our technique is somewhat a combination of the technique for the uniform-power physical interference model and the non-uniform disk model. Basically, we divide the base stations into power-classes depending upon their transmission power, and then, solve the allocation problem for each class independently. Finally, we pick the power-class that has the allocation with the highest social-welfare. The outline of our allocation algorithm is as follows.

1) Classify the base stations into $\lfloor \log(P_{\max}) \rfloor$ *power-classes*, based on the associated transmission power. In particular, power-class J contains base stations whose transmission power P_i lies in the range $P_i \in [2^J, 2^{J+1})$.
2) For each power-class J,
    a) Divide the network region into hexagons of side-length $R = (\sqrt[\alpha]{\beta} + 1)r/2$ each.
    b) Uniformly-color the hexagons using q_2 colors, where q_2 is as defined in Equation 3 later.
    c) Independently, for each hexagon H, allocate channels to base stations of power-class J contained in H. These set of base stations satisfy Property (P'1), and hence, this allocation subproblem is a multi-unit auction and we can use techniques described in Section III-A.
    d) For each color, combine the results from all hexagons of that color, since it can be shown that base stations of class J in different co-colored hexagons satisfy Property (P'2).
    e) Pick the color that has the highest total social-welfare.
3) The above gives an allocation for each power-class. Pick the allocation for the power-class that has the highest social-welfare.

We need to define q_2 as:

$$q_2 = \min\{x | x \geq q_2', x \geq 7, \text{ and } x = i^2 + j^2 + ij \text{ where } i, j \in \mathbb{Z}^+\}, \quad (3)$$

where q_2' is

$$q_2' = \left(\frac{4\sqrt{7}}{(3\sqrt{7} - 6)(\sqrt[\alpha]{\beta} + 1)} \right)^2 \left(\frac{12\beta}{(\alpha - 2)} \right)^{\frac{2}{\alpha}}.$$

Now, using arguments similar to the uniform-power physical interference and non-uniform disk model, the following result follows (we omit the proof here).

Theorem 4: For the TSA-MSW problem under the physical interference model with non-uniform transmission power, the auction mechanism based on the above described allocation algorithm and VCG payments is truthful and returns a valid spectrum allocation whose social-welfare is $2q_2$-approximate for the general-minded bidding model and is $q_2(1+\epsilon)$-approximate for the k-minded bidding model for a given $\epsilon > 0$. Here, q_2 is as defined in Equation 3. □

V. Simulation

In this section, we present our simulation results. The main purpose of our simulations is to demonstrate the efficiency of our designed auction mechanism in terms of multiple performance metrics. We start by describing our simulations set-up.

Network Topology and Model. In our simulations, we consider only unit-disk pairwise interference model due to space constraints and because for physical interference model, there is no simple truthful auction mechanism known that we could use for comparison. We consider two types of networks, as described below.

- *Random Networks:* We randomly place base stations within a fixed area of 1000×1000 square units. We vary the network density by varying the number of base stations from 50 to 1000 (with the default being 500). We use cells of uniform radius of 50 units.
- *Real Networks:* We use locations of real cellular base stations available in FCC public GIS database [31] and choose the 843 base stations deployed in the state of Massachusetts. Here, we choose a realistic cell radius of 10 kilometers.

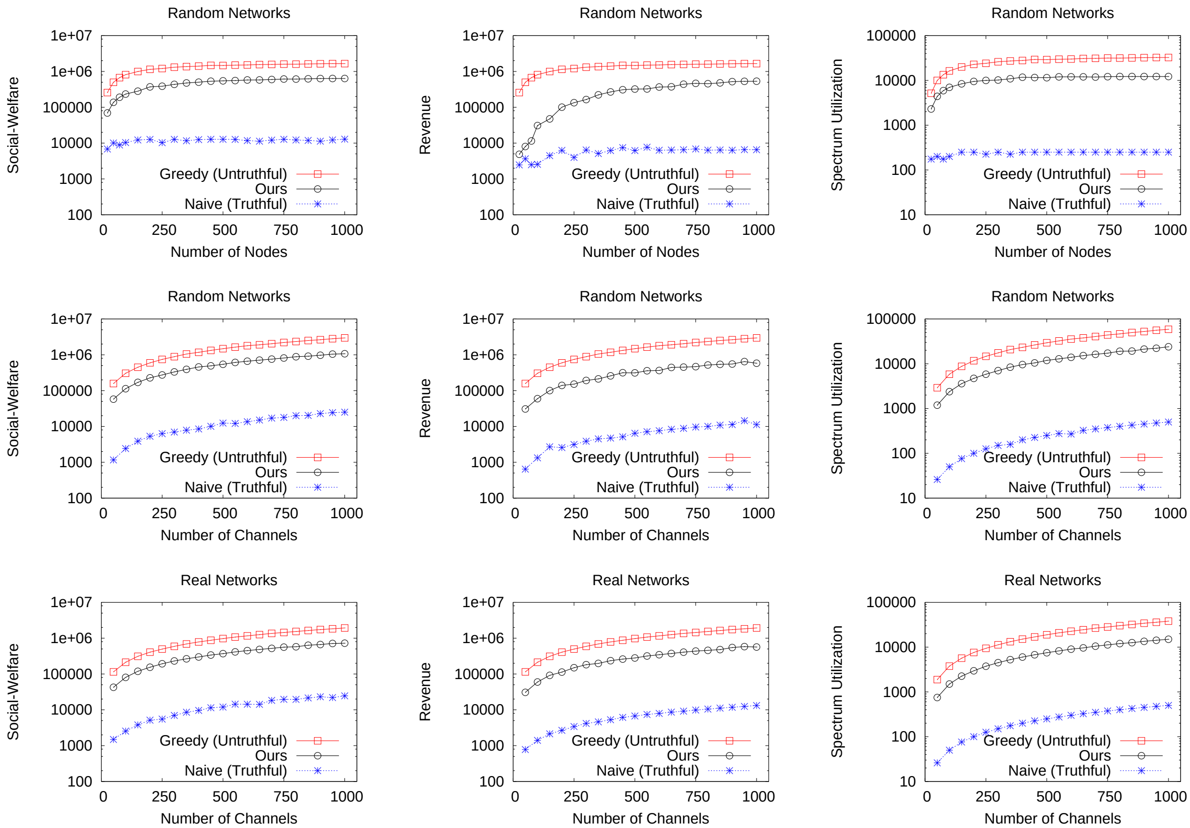

Fig. 5. Performance comparison of various auction mechanisms. The first six plots (in the first two rows) are for random networks with varying number of nodes (with 500 channels) and varying number of channels (with 500 nodes). The last three plots are for the cellular network in Massachusetts with 843 base stations and varying number of channels.

In both networks, we set up an auction of up to 1000 orthogonal single-type channels with the default being 500 channels; this is a reasonable range based on the past FCC auctions [17, 18].

Bidding Functions. We generate general-minded bidding functions for each base station i as follows. First, we randomly choose a non-zero *demand* for i, which is the maximum number of channels i is interested in. Then, we randomly generate i's bid for the first channel and "marginal" bids for each additional channel until the demand is satisfied. Beyond the demand, marginal bids for each additional channel is assigned zero (to satisfy the free-disposal property). Each marginal bid is chosen from the range $[0, 100]$. The above scheme results in valid general-minded bidding functions.

Auction Mechanisms Compared. We compare our auction mechanism with two auction mechanisms, viz., (i) Greedy, the best known (non-truthful) approximation spectrum allocation algorithm for maximizing social-welfare and/or revenue, and (ii) Naive, a simple truthful spectrum auction mechanism. Note that the only work on truthful spectrum auction mechanism is by Zhou et al. [14] which is restricted to simple single-minded or range bidding functions.

Greedy Auction Mechanism. Greedy is a non-truthful auction mechanism whose social-welfare as well as revenue is within a factor of 5 of the respective optimals, for the unit-disk model and non-complementary bidding function [11][8] – thus, it is the best and most general approximation algorithm known. Greedy's winner determination function allocates channels iteratively to the highest available bid without violating the interference constraint; this allocation results in a 5-approximate social-welfare [11]. If we charge each bidder a payment equal to its bid (declared valuation) for the allocated number channels, then Greedy's revenue is also within a factor of 5 of the optimal revenue possible.[9] Note that Greedy's social-welfare is equal to its revenue.

Naive Auction Mechanism. We now describe a simple auction mechanism (called Naive) that is truthful, but has no performance guarantee on the social-welfare and revenue. The Naive auction mechanism is closely based upon the Naive auction mechanism suggested by Zhou et al. [14] used as a comparison-benchmark in their work. Naive's allocation algorithm divides the entire network region into square grid of unit side-length.[10] Next, the algorithm uniformly colors the resulting square cells using 4 colors, and assigns each color $(1/4)^{th}$ of the available channels. This means that all square cells of the same color will use the same channels. Now, for each square cell H, Naive allocates all the channels usable in H to the bidder with the maximum bid for that many channels, and charges it a payment equal to the second highest bid in H. This is a simple generalization of Vickrey's auction [19] (see Section II) in each square cell. Finally, we note here that

[8]For complementary bidding function, the approximation ratio of Greedy is not guaranteed.

[9]For computing the optimal revenue, we assume that bidder's payment in an outcome must not be more than its declared valuation for the outcome.

[10]To ensure validity of the resulting allocation, the square cells are open from one side and closed from the other (similar to Figure 1).

Greedy is a pseudo-polynomial algorithm since its running time is polynomial in m, the number of channels, while our algorithm and Naive are polynomial in $\log(m)$.

Simulation Results. In our simulation, we compare Greedy, Naive, and Our (based on hexagonal division and coloring) auction mechanisms for the following three performance metrics: (i) social-welfare, (ii) revenue, and (iii) spectrum utilization. The *spectrum utilization* [14] is defined as the total number of allocation pairs in the spectrum allocation (see Definition 7). Spectrum utilization gives a measure of the spatial reuse of a spectrum allocation.

In Figure 5, we plot results for the above three metrics. For the random network, we vary number of base stations (nodes) as well as number of available channels, while for the fixed real network we vary the number of available channels. We observe that Greedy performs the best in all three performance metrics, but is only within a factor of 2 to 3 of that of our auction mechanism. Note that both Greedy and ours deliver an approximate social-welfare, and Greedy also delivers an approximate revenue, but is untruthful. Secondly, our auction mechanism outperforms the Naive mechanism by an order of magnitude, in all three performance metrics.

Thus, apart from the key properties of truthfulness and provably approximate social-welfare, our auction mechanism also delivers near-optimal revenue in practice. The simulation results also show that a Naive truthful auction mechanism can perform very badly.

VI. Conclusions

The recent trend of dynamic spectrum access in cellular networks creates a setting for auctioning of pieces of wireless spectrum to competing base stations. To mitigate market manipulation, a truthful spectrum auction is highly desired, so that bidders can simply bid their true valuations. For economic efficiency, we also want to allocate channels to bidders that value it the most. Thus, in this paper, we have designed a truthful spectrum auction that delivers an allocation with approximate social-welfare, for general interference and bidding models. Through simulations, we show that the revenue generated by our auction mechanism is also within a factor of 2-3 of the best-known approximation algorithm. In general, our mechanism performs an order of magnitude better than a Naive truthful spectrum auction mechanism.

References

[1] M. M. Buddhikot, P. Kolodzy, S. Miller, K. Ryan, and J. Evans, "DIMSUMnet: New directions in wireless networking using coordinated dynamic spectrum access," in *IEEE WoWMoM*, 2005.

[2] Spectrum Policy Task Force, "Spectrum Policy Task Force Report," *Federal Communications Commission ET Docket 02-135*, 2002.

[3] V. Brik, E. Rozner, S. Banerjee, and P. Bahl, "DSAP: A protocol for coordinated spectrum access," in *IEEE DySPAN*, 2005.

[4] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri, "A general framework for clearing auction of wireless spectrum," in *IEEE DySPAN*, 2007.

[5] S. Sengupta, M. Chatterjee, and S. Ganguly, "An economic framework for spectrum allocation and service pricing with competitive wireless service providers," in *IEEE DySPAN*, 2007.

[6] S. Talukdar, L. Lave, K. wei Lye, K. Marshall, and E. Subrahmanian, "Agents, evolutionary learning and market failure modes," in *36th Hawaii International Conference on Systems Science*, 2003.

[7] F. A. Wolak, "An empirical analysis of the impact of hedge contracts on bidding behavior in a competitive electricity market," *International Economic Journal*, vol. 14, no. 2, pp. 1–39, June 2000.

[8] M. Ilic, J. W. Black, E. Fumagalli, P. Visudhiphan, and J. L. Watz, "Understanding demand: The missing link in efficient electricity markets," Energy Laboratory Publication, Tech. Rep. MIT EL 01-014WP, 2001.

[9] C. Barrett, D. Cook, V. Faber, G. Hicks, A. Marathe, M. Marathe, A. Srinivasan, Y. J. Sussmann, and H. Thorn-quist, "Experimental analysis of bilateral-contracts," *Journal of Graph Algorithms and Applications*, vol. 7, no. 1, pp. 3–31, 2003.

[10] S. Borenstein, J. Bushnell, and F. Wolak, "Diagnosing market power in california's restructured wholesale electricity market," National Bureau of Economic Research, Inc, NBER Working Papers 7868, Sep. 2000.

[11] A. P. Subramanian, M. Al-Ayyoub, H. Gupta, S. R. Das, and M. M. Buddhikot, "Near-optimal dynamic spectrum allocation in cellular networks," in *IEEE DySPAN*, 2008.

[12] O. Ileri, D. Samardzija, T. Sizer, and N. B. Mandayam, "Demand responsive pricing and competitive spectrum allocation via a spectrum server," in *IEEE DySPAN*, 2005.

[13] Y. Wu, B. Wang, K. J. R. Liu, and T. C. Clancy, "A multi-winner cognitive spectrum auction framework with collusion-resistant mechanisms," in *IEEE DySPAN*, 2008.

[14] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "eBay in the sky: strategy-proof wireless spectrum auctions," in *MobiCom*, 2008.

[15] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2007.

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

[17] P. Cramton, *In Handbook of Telecommunications Economics*. Amsterdam: Elsevier, 2002, ch. Spectrum Auctions, pp. 605–639.

[18] A. Archer, C. Papadimitriou, K. Talwar, and E. Tardos, "An approximate truthful mechanism for combinatorial auctions with single parameter agents," *Internet Mathematics*, vol. 1, no. 2, pp. 129–150, 2003.

[19] W. Vickrey, "Counterspeculation, auctions and competitive sealed tenders," *Journal of Finance*, pp. 8–37, 1961.

[20] A. V. Goldberg and J. D. Hartline, "Collusion-resistant mechanisms for single-parameter agents," in *SODA*, 2005.

[21] E. H. Clarke, "Multipart pricing of public goods," *Public Choice*, pp. 17–33, 1971.

[22] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, pp. 617–631, 1973.

[23] D. Lehmann, L. I. Oćallaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," *J. ACM*, vol. 49, no. 5, pp. 577–602, 2002.

[24] A. Mu'alem and N. Nisan, "Truthful approximation mechanisms for restricted combinatorial auctions: extended abstract," in *Eighteenth national conference on Artificial intelligence*, 2002.

[25] X. Zhou and H. Zheng, "TRUST: A general framework for truthful double spectrum auctions," in *IEEE INFOCOM*, 2009.

[26] W. Stallings, *Data and Computer Communications*, 8th ed. Prentice Hall, 2007.

[27] S. Dobzinski and N. Nisan, "Mechanisms for multi-unit auctions," in *EC*, 2007.

[28] W. C. Y. Lee, *Mobile Cellular Telecommunications: Analog and Digital Systems*, 2nd ed. McGraw-Hill, 1995.

[29] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz, "Bundling equilibrium in combinatorial auctions," *Games and Economic Behavior*, vol. 47, no. 1, pp. 104–123, 2004.

[30] M. Al-Ayyoub and H. Gupta, "Joint routing, channel assignment, and scheduling under general interference models," under review.

[31] FCC Geographic Information Systems. [Online]. Available: http://wireless.fcc.gov/geographic/