Truthful Spectrum Auctions with Approximate Social-Welfare or Revenue

Mahmoud Al-Ayyoub,[1] Himanshu Gupta[2]
[1] Jordan University of Science and Technology, Jordan
[2] Stony Brook University, NY, U.S.A.

***Abstract*—In cellular networks, a recent trend in research is to make spectrum access dynamic in the spatial and temporal dimensions, for the sake of efficient utilization of spectrum. In one such 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. A key objective in designing an auction mechanism is "truthfulness." Combining this objective with an optimization of some social choice function (such as the social-welfare or the generated revenue) is highly-desirable. In this article, we design polynomial-time spectrum auction mechanisms that are truthful and yield an allocation with $O(1)$-approximate social-welfare or revenue. Our mechanisms generalize to general interference models. To the best of our knowledge, ours is the first work to design polynomial-time truthful spectrum auction mechanisms with a constant-factor approximation of either the expected revenue or the social-welfare. We demonstrate the performance of our designed mechanism through simulations.**

I. Introduction

Usage of wireless spectrum has long been governed by governmental regulatory authorities (e.g., FCC in USA or Ofcom in UK) who divide the spectrum into fixed size chunks to be used strictly for specific purposes, such as broadcast radio/TV, cellular/PCS services, wireless LAN, etc. This allocation is very long-term and space-time invariant, and is often based on peak usage. Such long-term allocation of spectrum introduces significant inefficiencies in utilization [1]. Thus, a new policy trend [2] is to make spectrum access dynamic. In case of cellular networks, centralized architectures [1, 3–5] for dynamic spectrum access have gained a lot of interest. In such models, a spectrum broker periodically 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 which made international headlines, leading to many studies [6–10].

A natural objective of an auction-based mechanism is to maximize the generated *social-welfare* (total "valuation" of the goods sold) or the generated *revenue* (total payments by the buyers) [4, 5, 11, 12]. However, a mechanism considering such an objective alone can encourage the spectrum buyers to lie about their real valuations (i.e., an "untruthful" auction), instill fear of market manipulation, and indirectly possibly lower revenue or social-welfare. Moreover, in a competitive environment, buyers may spend a lot of time/effort in predicting the behavior of other buyers and planning against them. In this article, our focus is on designing spectrum auction mechanisms that not only encourage truthful behavior but also provide some form of guarantee on either the social-welfare or the revenue.

Model and Contributions. In a spectrum auction, the items being sold are various channels corresponding to certain blocks of frequency. The base stations bid for these channels, based on their valuations. The auctioneer assigns channels to base stations within the "wireless interference constraint" and determines payments from bidders. In the above context, we wish to design polynomial-time auction mechanisms that (i) encourage buyers to be truthful (i.e., ensure that the buyers "benefit" the most when their bid is equal to their actual valuation), and (ii) maximize either the social-welfare or the revenue. However, when the bidder valuations are completely private, no truthful auction mechanism can give any performance guarantee on the revenue (see Section IV). Keeping the above in mind, we make the following contributions in the paper:

- When the bidder valuations are private, we design truthful spectrum auctions with approximate social-welfare.
- We design truthful spectrum auctions with approximate revenue, under the relaxed Bayesian setting wherein the bidder valuations are drawn from publicly-known probability distributions.
- We extend our above designed mechanisms to general interference models and other generalizations.

To the best of our knowledge, ours is the first work to design polynomial-time truthful spectrum auctions that offer a constant-factor approximation on the social-welfare or the expected revenue.

Paper Organization. We start our discussion with some background material on spectrum auctions (Section II). In Section III, we design truthful mechanisms with near-optimal social-welfare, and in Section IV, we consider the Bayesian setting and design a truthful mechanism with near-optimal expected revenue. Finally, we compare the performance of our proposed mechanisms with known works in Section V.

II. Background and Related Works

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.

Dynamic Spectrum Access. In the dynamic spectrum access architectures, the spectrum is allocated dynamically in spatial and temporal domains, 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). The base stations bid for these available channels by specifying a bidding price. Periodically, the spectrum broker allocates the channels to the base stations under the "wireless interference constraint" such that the total revenue (sum of payments 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. However, 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 (truthful) auction mechanisms. We also discuss the most general example of truthful auction mechanisms.

Auction Mechanism. In an auction [13], 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 $\mathbf{v}_i : O \mapsto \mathbb{R}$ called its *valuation function*, which quantifies the bidder's benefit from each outcome.
- Bidders are asked to declare their valuation functions;[1] let $\mathbf{w}_i$ denote the *declared valuation function* of the i^{th} bidder. The bidders may lie about their valuation functions; thus $\mathbf{w}_i$ may not be equal to $\mathbf{v}_i$.
- 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 payment p_i.
- *Utility* u_i of a bidder i is the difference between its true valuation of the chosen outcome o and its payment p_i, i.e., $u_i = \mathbf{v}_i(o) - p_i$. Each bidder's goal is to maximize its utility.

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

- The winner determination function $\mathbf{x}$ accepts as input a vector $\mathbf{w} = (\mathbf{w}_1, \ldots, \mathbf{w}_n)$ of bidding (declared valuation) functions and returns an output $\mathbf{x}(\mathbf{w}) \in O$.
- The payment function $p(\mathbf{w}) = (p_1(\mathbf{w}), \ldots, p_n(\mathbf{w}))$ returns a real vector quantifying the payment charged by the mechanism to each of the bidders. □

[1] Such auctions (wherein bidders declare their valuations) are called *direct revelation* auctions. Though more general types of auctions exist, in this paper we focus only on direct revelation auctions.

Definition 2: (**Social-Welfare; Revenue**) The *social-welfare* of an outcome o is defined as the sum of the valuations, i.e., $\sum_i \mathbf{v}_i(o)$. On the other hand, the *revenue* of an auction mechanism $(\mathbf{x}, p)$ is the sum of the payments $\sum_i p_i(\mathbf{w})$ charged to the bidders for a given declared valuation vector $\mathbf{w}$. □

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* mechanism enforces bidders to behave truthfully by offering them incentives in the form of reduced payments. These incentives are based on the presumption that each bidder's objective is to maximize 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's utility is maximized when it truthfully declares its valuation function $\mathbf{v}_i$.

More formally, let the true valuation functions of the bidders be $\mathbf{v} = (\mathbf{v}_1, \ldots, \mathbf{v}_n)$. Consider two declared valuation function vectors, viz., (i) $\mathbf{w} = (\mathbf{w}_1, \ldots, \mathbf{w}_{i-1}, \mathbf{v}_i, \mathbf{w}_{i+1}, \ldots, \mathbf{w}_n)$, and (ii) $\mathbf{w}' = (\mathbf{w}_1, \ldots, \mathbf{w}_{i-1}, \mathbf{w}_i, \mathbf{w}_{i+1}, \ldots, \mathbf{w}_n)$ (where $\mathbf{w}_i \neq \mathbf{v}_i$). A mechanism $(\mathbf{x}, p)$ is considered *truthful* if $\mathbf{v}_i(\mathbf{x}(\mathbf{w})) - p_i(\mathbf{w}) \geq \mathbf{v}_i(\mathbf{x}(\mathbf{w}')) - p_i(\mathbf{w}')$ for all $\mathbf{v}$, i, and $\mathbf{w}_i$. □

VCG Mechanisms. The only general mechanism that guarantees truthfulness is due to Vickrey-Clarke-Groves (VCG) [14–16]. 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 others with and without i's participation.

Definition 4: (VCG Mechanism.) A VCG mechanism is an auction mechanism $(\mathbf{x}, p)$ that satisfies the following two conditions, for any given declared valuation functions $\mathbf{w} = (\mathbf{w}_1, \ldots, \mathbf{w}_n)$.

- $\mathbf{x}(\mathbf{w}) \in \arg\max_o \sum_i \mathbf{w}_i(o)$, i.e., the winner determination function $\mathbf{x}$ chooses an outcome that maximizes the social-welfare according to $\mathbf{w}$.
- The payment functions are determined by the VCG formula $p_i(\mathbf{w}) = (-\sum_{j \neq i} \mathbf{w}_j(\mathbf{x}(\mathbf{w}))) + h_i(\mathbf{w}_{-i})$, where each $h_i(\mathbf{w}_{-i})$ is an arbitrary function of $\mathbf{w}_{-i} = (\mathbf{w}_1, \ldots, \mathbf{w}_{i-1}, \mathbf{w}_{i+1}, \ldots, \mathbf{w}_n)$.

□

It can be shown that VCG mechanisms maximize social-welfare. In fact, VCG mechanisms are the only general truthful mechanisms with optimal social-welfare allocations. However, they require solving an optimization problem which can be NP-hard in many settings, as is the case in our context of spectrum auctions. Furthermore, they may result in low (even zero) revenue in some cases.

B. Related Works on Spectrum Auctions

Truthful Spectrum Auctions. To the best of our knowledge, there has been only four works till date, viz., [17–20], that have designed truthful mechanisms for spectrum auction. Below, we discuss each one of them.

The truthful mechanism designed by Zhou et al. [17] does not attempt to maximize the revenue or social-welfare. Moreover, their approach is limited to only simple pairwise interference model. As observed in [17, 21], it is rather straightforward to design a truthful auction mechanism without any regard for revenue or social-welfare. However, the authors in [17] do show through simulations that their mechanism returns better revenue and social-welfare compared to a simple truthful mechanism. Recently, this work has been extended to consider double auctions [22].

In a recent work, Jia et al. [18] present a simple extension of Myerson's mechanism [23] for spectrum auctions. However, the extension results in an exponential-time mechanism, since the corresponding virtual-surplus maximizing problem is NP-hard. Realizing the seriousness of this shortcoming, [18] presents a polynomial-time mechanism based on the greedy mechanism of [24]. However, the expected revenue delivered by such a mechanism can be arbitrarily bad, as shown in Section IV-B.

In another work, Wu et al. [19] design a spectrum auction mechanism based on the truthful VCG mechanism [13]. 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 base stations [25]. Finally, they assume either a single-channel system or that each bidder is interested in only one channel in a multi-channel system.

Finally, in a recent concurrent and complementary work, Hoefer et al. [20] design approximation algorithms for allocation of channels to wireless nodes to maximize the social welfare (without worrying about truthfulness), when nodes bid for sets of channels. They suggest combining their technique with the framework of [26] to obtain *randomized* mechanisms that are truthful in *expectation*. They do not consider maximization of revenue. Also, we note that the physical interference model considered in [20] is unsuitable for cellular networks,[2] which is our main focus.

Other Works. Recently there have been lots of works on dynamic spectrum allocation [4, 5, 11, 12, 27] using either auction-based or pricing-based mechanisms, but all of these works have ignored the truthfulness property.

III. Social-Welfare Maximization

In this section, we design a truthful spectrum auction with near-optimal social-welfare. The objective of maximizing social-welfare is justified in many settings and is extensively studied in economics [13, 28]. In particular, a spectrum auction mechanism with high social-welfare ensures that the spectrum is sold to bidders who value it the most and hence, are likely to utilize it most efficiently.

As mentioned before, VCG mechanism for spectrum auctions requires solving an NP-complete problem. Solving the NP-hard optimization problem using an approximation algorithm is not helpful, since it destroys the truthfulness of the mechanism [13]. To circumvent this, the authors of [13] develop the below described maximal-in-range (MIR) technique, which can be used to design polynomial-time truthful mechanisms with near-optimal social-welfare in certain settings.

Maximal-In-Range Mechanisms. In [13], the authors 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 $\mathbf{x}$ at V, i.e., $\mathcal{O} = \{\mathbf{x}(\mathbf{v}) | \mathbf{v} \in V\}$. We say that $\mathbf{x}$ is *maximal in its range* if for every $\mathbf{v} \in V$, $\mathbf{x}(\mathbf{v})$ maximizes the social-welfare over $\mathcal{O}$. □

<u>MIR for Multi-Unit Auctions.</u> In a multi-unit auction (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 [29], the authors design an MIR mechanism for multi-unit auctions that is truthful and yields an allocation with approximate social-welfare. We use their technique for our auction mechanism design, as described below.

Our Approach. 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 auctions from [29] 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 ratio of the social-welfare.

A. Truthful Spectrum Auction with Approximate Social-Welfare

Throughout the article, we use the terms "bidder" and "base station" interchangeably.

Spectrum Auction Model. 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 a certain number of channels.

[2] In cellular networks, for a base station x, the intended receivers are *any* mobile devices that are within a certain distance of x (as in Section III-B).

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

Representation of Valuation and Bidding Functions. In a spectrum auction of channels of the 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 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 bidder i is a declaration of its privately-known valuation function v_i. Thus, the bidding function 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 are 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.

Interference Graph. Each base station is associated with a region around it called its *coverage-cell*; each base station serves its clients in its coverage-cell. To communicate, the base station and the client must operate "interference-free" on a channel. In cellular networks, wireless interference at a client may arise due to multiple near-by base stations operating on the same channel. In a simple model of *pairwise interference*, pairs of base stations with intersecting coverage-cells are said to *interfere* with each other if operating on the same channels, and thus, must not be assigned a common channel. Such a relationship between pairs of base stations can be represented by edges 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 coverage-cells of the corresponding base stations intersect. □

If the coverage-cells of the base stations are unit-radius disks, then the interference graph is a *unit-disk graph*. ***For clarity of presentation,*** *we assume unit-disk interference graph below.* Our technique can be easily generalized to more involved pairwise interference models [30]; in Section III-B, we present the generalization of our technique to the physical interference model.

Valid Spectrum Allocation. Given an interference graph, the spectrum allocation must be done in such a way that no pair of interfering base stations are allocated a common channel. This *interference constraint* is incorporated in the below definition of a valid spectrum allocation.

Definition 7: (Valid Spectrum Allocation.) Let V and C be the set of base stations and available channels, and let $P(C)$ denote the power set of C. A spectrum allocation vector $(x_1, \ldots, x_{|V|})$ is considered *valid* if there is an assignment $a : V \mapsto P(C)$ such that (i) $|a(i)| \geq x_i$ for all i, and (ii) $a(i) \cap a(j) = \emptyset$ if (i, j) is in E_t. □

It can be shown it is NP-complete to test whether a given allocation vector is valid, through a reduction from the problem of partitioning a graph into minimum number of independent sets. Thus, it is desirable for the auction mechanism to output the assignment function a *in addition (or in lieu of)* to the allocation vector.

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.

TSA-MSW problem is NP-hard even without the truthfulness objective [11]. Thus, we focus on designing a truthful mechanism that yields approximate social-welfare.

Truthful Mechanism with Approximate Social-Welfare. Given a network with base stations, the unit-disk interference graph, and the bidding functions, we first determine a valid allocation with approximate social-welfare as follows. Basically, we divide the entire network into small hexagonal regions, solve the simpler optimization problem in each hexagon independently, and then, "combine" the solutions. At a high-level, our algorithm consists of the follows steps.

1) Divide the entire network region into small hexagons of unit side-length.
2) Uniformly-color the hexagons with seven colors. See Figure 1.
3) Allocate channels to base stations in each hexagon *independently*, treating it as a multi-unit auction (MUA) and using techniques similar to [29]. 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 and allocate the channels to the winners accordingly.
6) Charge the winners VCG payments (see Definition 4).

Properties of Coloring. In the coloring suggested in Step 2 above, the following two properties hold.

`Property 1`	Every pair of base stations in the same hexagon interfere with each other.
`Property 2`	Base stations in different co-colored hexagons do not interfere with each other.

`Property 1` follows directly from the definition of unit-disk interference, while `Property 2` follows from the fact that the distance between base stations in different co-colored hexagons will be at least $(\sqrt{3(7)} - 2) > 2$ (from Lemma 2).

Details of Step 3. 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. Thus, we use [29]'s techniques within each hexagon as described below for general-minded and k-minded bidding functions.

General-Minded Bidding Function. In case of general-minded bidding model, the available m channels are split into

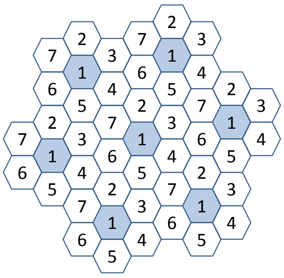

Fig. 1. Hexagons uniformly-colored using 7 colors.

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 polynomial time. This yields an allocation with 1/2-approximate social-welfare [29] within a hexagon.

k-minded Bidding Function. For k-minded bidding functions, a restricted form of allocation known as the *t-round* allocation is used. For a given t (where t is a PTAS parameter), 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 using dynamic programming. Finally, the best allocation among the $tN_H^t k^t$ such allocations is picked as the optimal t-round allocation. The above allocation algorithm runs in polynomial time for a fixed t, and yields a $\left(1 - \frac{1}{t+1}\right)$-approximate allocation [29] within a hexagon.

Proof of Truthfulness and Approximation.

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 $\epsilon > 0$. Also, the mechanism's running time is polynomial in n and $\log m$, where n is the number of nodes and m is the number of channels.

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 [13].

Approximate Social-Welfare. First, note that by `Property 1` and `Property 2` 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 $\mathcal{A}_c$ be the allocation constructed by our algorithm and $\mathcal{O}_c$ be the allocation with optimal social-welfare. We show that the social-welfare of $\mathcal{A}_c$ is within a factor of 2 of that of $\mathcal{O}_c$. 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 $\mathcal{A}_c$'s ($\mathcal{O}_c$'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 $\mathcal{A}_c$ is within a factor of 2 of that of $\mathcal{A}_c$. 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. ∎

Pseudo-polynomial Algorithms. Note that instead of the above allocation algorithms from [29] within a hexagon, we could also use the optimal dynamic programming approach which runs in $O(m^4 N_H)$ time. Since the size of the input is $O(\log m)$, this optimal dynamic programming algorithm has a pseudo-polynomial time complexity. However, in our simulations, we observe that this pseudo-polynomial algorithm does not perform any better than our polynomial algorithms.

B. Extensions

Our technique can be easily generalized to more involved pairwise interference models and non-orthogonal channels, as shown in [30]. In this subsection, we present the generalization of our technique to the physical interference model.

Physical Interference Model. In the physical interference model, a reception from a base station i is successful at a point p if and only if,

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

where P is the uniform transmission power, B' is the set of other base stations operating on the same channel as i, δ_x is the distance of the point p from a base station x, $\mathcal{N}$ is the ambient noise, and α is the path loss exponent.

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.

Valid Spectrum Allocation. In the physical interference model, a spectrum allocation vector $(x_1, \ldots, x_n)$ is considered valid if there is an assignment function $a : V \mapsto P(C)$ such that (i) $|a(i)| \geq x_i$ for all i, and (ii) for any i and c such that $c \in a(i)$, the SINR of channel c at any point p within a distance of r from i should be greater than β, i.e., $(P/\delta_i^\alpha)/(\mathcal{N} + \sum_{j \in B} P/\delta_j^\alpha) \geq \beta$ where B is the set of base stations j such that $c \in a(j)$ and δ_x is the distance of x from p.

Hexagonal Division and Coloring. Here, we need to do the hexagonal division and coloring in a way so as to satisfy the following two properties: `Property 3:` No two base stations in a hexagon can be allocated the same channel, and (ii) `Property 4`: If one node from each hexagon of the same color is concurrently active on the same channel, then transmission from these nodes is successful within their communication radius.

Plane Division and Coloring. It is easy to see from the SINR Equation 1 that dividing the network region into hexagons of side-length

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

would ensure `Property 3`. Here (and in Lemma 1 below), for simplicity, we have assumed the ambient noise $\mathcal{N}$ to be zero; nonzero noise can be incorporated using techniques similar to [31]. Now, to determine appropriate coloring needed to satisfy `Property 4`, we state the following three lemmas.

Lemma 1: Given a division of the region into hexagons of side-length r', `Property 4` is satisfied if the minimum distance between co-colored hexagons is at least $\sqrt{3q_1'}r'$, where q_1' is

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

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[4] of at least $(\sqrt{3q_1}-2)r'$ from H (from Lemma 2). 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:

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

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; nonzero noise can be incorporated using techniques similar to [31]. Now, using the value of q_1 from Equation 4, the SINR at point p due to the transmission at base station i is at least:

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

■

The following two lemmas are derived from [25, 32].

Lemma 2: In a hexagonal division with side-length r' and uniformly-colored with x colors, the distance between the centers of co-colored hexagons is at least $\sqrt{3x}r'$. ■

Lemma 3: 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. ■

The below theorem follows from the above three lemmas.

Theorem 2: Given a division of the region into hexagons of side-length r', the number of colors q_1 required to satisfy `Property 4` is given[5] by:

$$q_1 = \min\{x \mid x \geq \max(7, q_1'), \text{ and } x = i^2+j^2+ij \text{ where } i,j \in \mathbb{Z}^+\}. \qquad (4)$$

■

Overall Allocation Algorithm. As discussed above, dividing the region into hexagons of side-length r' (Equation 2) and coloring them uniformly using q_1 (Equation 4) colors, allows us to satisfy `Property 3` and `Property 4`. `Property 3` ensures that allocation in each hexagon can be treated as a multi-unit auction, while `Property 4` allows us to re-use channels across hexagons of the same color. Now, using the same allocation algorithm as before, we have the following.

Theorem 3: For the TSA-MSW problem under the physical interference model, the above described mechanism 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 $\epsilon > 0$, where q_1 is as defined in Equation 4. ■

IV. Revenue Maximization under the Bayesian Setting

We now design a truthful spectrum auction mechanism with near-optimal revenue. In an auction setting, wherein the bidder's valuation is private, it is impossible to design truthful auctions with any guarantee on revenue [21, 33]. To circumvent this, researchers have considered the Bayesian setting wherein each bidder's valuation is drawn from a known probability distribution [28]. In a seminal work [23], Myerson presents a truthful optimal mechanism for a single-item auction under the Bayesian setting. In this section, we essentially extend this classical result to spectrum auctions.

A. Truthful Auction Mechanisms under the Bayesian Setting

In this subsection, we briefly present the key points [28] of Myerson's mechanism applied to the more general single-parameter auctions. We start with basic definitions.

Single-parameter Auctions. In a *single-parameter* auction, each bidder i has a publicly-known set of outcomes $O_i \subset O$

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

[5]Note that in our context we should use at least 7 colors, irrespective of the values of α and β.

known as its *winning alternatives* and a private *valuation-value* v_i (a single value) such that $\mathbf{v}(o) = v_i$ for every $o \in O_i$ and $\mathbf{v}(o) = 0$ for every $o \notin O_i$. Bidders declare (perhaps, untruthfully) their valuation-value as their *bid* w_i.

(Valid) Allocation Vector. In a single-parameter auction, an outcome can be represented by an *allocation vector* of n binary variables $x = (x_1, \ldots, x_n)$, where x_i is 1 if the bidder i wins and zero otherwise. However, not all 0-1 vectors of length n may correspond to an outcome of the mechanism. The 0-1 vectors that correspond to an outcome are referred to as *valid allocation vectors*. For instance, in a single-item auction with 4 bidders, wherein the item is given to one of the 4 bidders, (0,0,0,1) is a valid allocation vector while (0,1,1,0) is not a valid allocation vector.

Myerson's Optimal Mechanism. Given, for each bidder i, the winning alternatives O_i, declared valuation-value (bid) w_i, and the distribution F_i of the private valuation-value v_i, the mechanism finds an allocation vector and payments such that truthfulness is maintained and the expected revenue is optimal where the expectation is taken over the randomness in bidders' valuations [28]. Myerson's mechanism is based on the following characterization of truthful mechanisms for single-parameter auctions.

Theorem 4 ([28, Theorem 13.6]): Consider a single-parameter auction, wherein the losers pay nothing (i.e., $x_i = 0$ implies $p_i = 0$). Under the Bayesian setting, a mechanism is truthful if and only if, for any bidder i and any fixed choice of bids by the other bidders:

(i) x_i is monotonic nondecreasing in w_i, and
(ii) the payment p_i for any winning bidder i is set to the *critical value* t_i, which is the minimum value i needs to bid in order to win. Note that, in general, t_i depends upon the bids of the other bidders. ∎

Given the above theorem, to specify a truthful mechanism, we need to only specify a winner determination function that satisfies the first condition of the theorem; the payments can be derived from the second condition. In [23], Myerson specifies the winner-determination function based on "virtual-bids," and shows that it leads to optimal expected revenue, if the payments are determined as described above.

Virtual Bids and Surplus. Myerson's mechanism [28] starts by replacing each bid w_i with a *virtual-bid* $\phi_i(w_i)$ as follows.

$$\phi_i(w_i) = w_i - \frac{1 - F_i(w_i)}{f_i(w_i)}, \tag{5}$$

where $f_i(x) = \frac{d}{dx}F_i(x)$ is the probability density function.

For a given outcome $o = (x_1, x_2, \ldots, x_n)$, the *virtual surplus* is defined as the sum of winning virtual-bids, i.e., $\sum_i x_i \phi_i(w_i)$. The following theorem is key to the design of an optimal truthful mechanism.

Theorem 5 ([28, Theorem 13.10]): The expected revenue of any truthful mechanism under the Bayesian setting is equal to its expected virtual surplus. Here, the expectations are taken over the distributions of the valuations. ∎

Myerson's Mechanism, and its Extensions. Myerson's mechanism essentially determines an outcome that maximizes the virtual surplus, and uses payments based on condition (ii) of Theorem 4. By the virtue of the above two theorems, such a mechanism will be truthful and optimal, if (and only if) the $\phi_i(w_i)$'s are monotonically nondecreasing in w_i [28].

Myerson's technique can be easily extended to more general single-parameter auctions [34–36]. Some other works have also extended Myerson's technique to simple multi-parameter settings [37–39].

B. Truthful Spectrum Auction with Approximate Expected Revenue

We consider the same model as before, except for the bidding function described below.

Bidding Function. Each base station (bidder) i has a publicly-known *demand* for d_i number of channels. Any outcome wherein i gets at least d_i channels is a *winning alternative* for i. Each bidder i also declares its valuation-value (bid) w_i for the winning alternatives, which may be different than its private valuation-value v_i. Note that the outcomes wherein i gets less than d_i channels are valued at zero by i. Here, we consider the *Bayesian setting*, wherein the valuation-value v_i is drawn randomly from a publicly-known probability distribution F_i. The above bidding model is much simpler than the bidding functions considered in previous section, since Myerson's technique is limited to the simple single-parameter setting. In Section IV-C, we generalize our technique to certain more general bidding functions.

TSA-MER (Truthful Spectrum Auctions with Maximum Expected Revenue) Problem. Given an interference graph, the number of available channels, and the bid-demand pair of each base station along with the distribution from which the valuation was drawn, the *TSA-MER problem* is to design a truthful auction mechanism that returns a valid spectrum allocation with maximum expected revenue.

The TSA-MER problem can be shown to be NP-hard, by a reduction from the maximum independent set problem, since maximizing expected revenue is equivalent to maximizing virtual surplus (sum of virtual-bids).

Recent Work on TSA-MER. In a recent work, Jia et al. [18] extended Myerson's mechanism for the TSA-MER problem.[6] However, since maximization of virtual surplus is NP-hard due to the interference constraint, Myerson's technique only yields an exponential-time mechanism. Thus, [18] designed a Greedy-heuristic mechanism for the TSA-MER problem, which considers nodes in decreasing order of virtual-bid per channel (i.e., $\phi_i(w_i)/d_i$) and allocates channels to them if the interference constraint is not violated. To check the interference constraint efficiently, we need to maintain the channels-to-bidders assignment function. Finally, the payments by the winners are determined as suggested in Theorem 4. The above mechanism is truthful, but the revenue yielded can be arbitrarily bad. (see Figure 2).

[6] We note that [18] actually considers a more general model wherein bid-demand pairs are associated with a service provider which controls multiple base stations. We consider such a generalization in Section IV-C2.

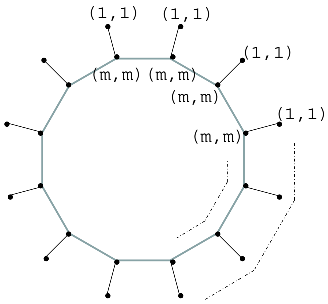

Fig. 2. Counter example for the Greedy mechanism. The figure shows the interference graph over given base stations. The (demand, bid) pair for the "inner" bidders is (m, m), while for the "outer" bidders is $(1, 1)$; here m is the total number of channels. The bids are constant, and hence, virtual-bid of each bidder is equal to its bid. Since all the bidders have the same rank (= virtual-bid/demand), the Greedy mechanism may pick all the outer bidders and yield a total revenue of $m/2$, while the optimal revenue is $m^2/4$.

Outline of the Truthful Mechanism with Approximate Expected Revenue. Based on Theorems 4 and 5 of Section IV-A, our method for designing a truthful spectrum auction mechanism with approximate expected revenue is outlined in the following two steps:

1) Determine a valid spectrum allocation with approximate virtual surplus, satisfying condition (i) of Theorem 4.
2) Determine payments using condition (ii) of Theorem 4.

We discuss the above steps in the following paragraphs.

Valid Allocation with Approximate Virtual Surplus. Given a network with base stations, the unit-disk interference graph, the demand-bid pairs, and the probability distributions of the bidder valuations, we determine a valid allocation with approximate virtual surplus using a technique similar to the one used in Section III-A except for the following changes. First, we have a pre-processing step in which we replace each bid w_i with a *virtual-bid* $\phi_i(w_i)$ as defined by Equation 5. Second, the problem of optimally allocating channels to base stations within each hexagon is treated as a Knapsack problem (rather than a MUA) where the virtual-bids are the "values" of items to be placed in the knapsack and the demands are their "weights." The well-known fully polynomial-time approximation scheme (FPTAS) [40] can be used to get a $(1 + \epsilon)$-approximate virtual surplus of each hexagon for any $\epsilon > 0$. Third, we pick the color that has the highest total virtual surplus (rather than social-welfare). And finally, we perform a post-processing step to greedily satisfy the demands of more base stations and charge the winners payments computed using condition (ii) of Theorem 4, as described below.

Post-Processing Step. We will show in Theorem 6 that the above allocation algorithm satisfies the monotonicity of x_i's (i.e., the first condition of Theorem 4). Incidentally, we can further improve the above allocation algorithm, *without violating the monotonicity of x_i's* (as will be shown in Theorem 6), by allocating more bidders in a greedy manner. In particular, we sort the *remaining* bidders by their virtual-bids per demand (i.e., $\phi_i(w_i)/d_i$), and consider them for allocation in that order without violating the interference constraint. To efficiently implement the above, we would need to maintain the channels-to-bidders assignment function. We note that the above post-processing however does not improve the approximation factor of our algorithm.

Determining Payments. The payments are determined according to Theorem 4 as follows. For each winner i, we use a binary search to find its critical value t_i (for the given fixed bids of other bidders) such that i wins if $w_i \geq t_i$ and loses otherwise. Note that such a value t_i is guaranteed to exist, since our allocation algorithm results in monotonically nondecreasing x_i's. Then, for each such winning bidder, we set its payment p_i as t_i. Losing bidders pay zero.

The critical values for bidders who win in the post-processing step can be determined using ideas based on the "critical neighbor" technique of [18]. The critical value for a bidder i who wins in the first step (involving coloring of hexagon cells) can be computed using at most $\log w_{\max}$ runs of the allocation within its hexagon cell[7] followed by the above "critical neighbor" technique; here w_{max} is the maximum valuation-value of any bidder. The latter part may be needed to determine the critical value for i's win due to the post-processing step; note that even if lowering the bid of i makes its hexagon color a loser in the first step, bidder i can still win due to the post-processing step.

Proof of Truthfulness and Approximation.

Theorem 6: For the TSA-MER problem under the Bayesian setting and the pairwise interference with unit-disk model, the above described mechanism is truthful and returns a valid spectrum allocation whose expected revenue is at least $\frac{1}{7(1+\epsilon)}$ of the optimal expected revenue, for a given $\epsilon > 0$. Also, the mechanism runs in time polynomial in $1/\epsilon$, n, and $\log m$.

Proof: The approximation proofs follow from the same arguments as in the proof of Theorem 1. Truthfulness. By Theorem 4, we need to only show that our allocation algorithm results in monotonically nondecreasing x_i's. First, note that the FPTAS algorithm used in each hexagon is monotonic since the FPTAS algorithm is an optimal algorithm over "scaled-down" values and the optimal algorithm is trivially monotonic. Now, to show the monotonicity of our overall mechanism, we need to consider two cases: (i) when a bidder i is selected as a winner in the first step, and (ii) when a bidder i is selected as a winner in the post-processing step. In the first case, if the bids of all other bidders remain fixed, then an increase in the bid of i would not change (a) the presence of i in the FPTAS knapsack-solution (due to its monotonicity), and (b) the winning of the color of i's hexagon. In the second case, increasing the bid of i will maintain its inclusion in the greedy post-processing step until the color of i's hexagon becomes a winning color. However, when the color of i's hexagon becomes a winning color (due to the increase in i's bid), i must still remain a winner in its hexagon (otherwise its hexagon's color would not have become a winning color).

Valid Spectrum Allocation. By virtue of `Property 2` and the fact that the allocation within each hexagon is a Knapsack solution, the allocation constructed before the post-processing step is valid. Since the post-processing step does not violate the interference constraints, the spectrum allocation returned by the designed mechanism is valid. ∎

[7] Note that the allocation within other hexagon cells does not change with the variation in i's bid.

C. More General Bidding Functions

The above mechanism can be extended to more involved pairwise interference and physical interference models using techniques similar to [30] and Section III-B. Here, we discuss generalization to more general bidding functions.

1) **Beyond Single-Minded Bidding:** We now extend our technique beyond single-minded bidding by handling *fractional demands*. More formally, a bidder i's declared demand-bid is of the form $(\check{d}_i, \hat{d}_i, w_i)$, signifying that the bidder would accept any number of channels between $\check{d}_i$ and $\hat{d}_i$ at a price of at most w_i per channel. For simplicity, we first assume that $\check{d}_i = 0$ for every bidder i; we relax this assumption later.

For the above setting, the mechanism's output is an allocation vector $(x_1, \ldots, x_n)$ wherein $x_i \in [0, 1]$ represents the fraction of demand satisfied, i.e., for a given x_i, the number of channels allocated is $x_i \hat{d}_i$. Also, for a given allocation vector, the virtual surplus is defined as $\sum_i \phi_i(w_i)\hat{d}_i x_i$. For this setting of fractional demands, Theorems 4 and 5 can be generalized (based on [28]) as follows. Below, we use the notation $x_i(w_i)$ to denote x_i for a given w_i and fixed bids of other bidders.

Theorem 7: A mechanism (wherein losing bidders pay zero, i.e., $x_i = 0$ implies $p_i = 0$) is truthful iff for any bidder i and any fixed choice of bids by other bidders,

- $x_i(w_i)$ is monotonically nondecreasing in w_i.
- The payment is set as follows

$$p_i(w_i) = w_i \hat{d}_i x_i(w_i) - \int_0^{w_i} \hat{d}_i x_i(t) dt. \qquad (6)$$

Proof: For simplicity, we drop the subscript i. To show truthfulness, we only need to show that the utility of truthful bidding v is no smaller than bidding any other value w. I.e.,

$$\begin{aligned} v\hat{d}x(v) - p(v) &\geq v\hat{d}x(w) - p(w) \\ \int_0^v \hat{d}x(t)dt &\geq v\hat{d}x(w) - w\hat{d}x(w) + \int_0^w \hat{d}x(t)dt. \end{aligned}$$

For $w > v$, the above is true since $(w - v)(\hat{d}x(w)) \geq \int_v^w \hat{d}x(t)dt$ follows from the monotonicity of x, while for $w < v$, the above is true since $(v - w)(\hat{d}x(w)) \leq \int_w^v \hat{d}x(t)dt$ also follows from the monotonicity of x.

Now to show the other direction, we take the truthfulness constraints at v, $v\hat{d}x(v) - p(v) \geq v\hat{d}x(w) - p(w)$, and at w, $w\hat{d}x(v) - p(v) \leq w\hat{d}x(w) - p(w)$. Rearranging these inequalities gives

$$v\hat{d}(x(w) - x(v)) \leq p(w) - p(v) \leq w\hat{d}(x(w) - x(v)). \qquad (7)$$

From this, we get $(w - v)(x(w) - x(v)) \geq 0$ which implies the monotonicity of x.

We now derive Equation 6. Let $w = v + \epsilon$, then, by dividing Equation 7 by ϵ and taking the limit, we get

$$v\hat{d}\ \frac{dx}{dv} \leq \frac{dp}{dv} \leq v\hat{d}\ \frac{dx}{dv}.$$

Now, since $p(w) = 0$ for any w smaller than the critical value, we get

$$p(w) = \int_0^w t\hat{d}x'(t)dt.$$

Integrating the above equation by parts gives Equation 6. ∎

Theorem 8: Under the Bayesian setting with fractional demands, the expected revenue of any truthful mechanism is equal to its expected virtual surplus, where the virtual surplus is as defined above. ∎

Overall Mechanism. On the basis of the above two theorems, our mechanism from the previous subsection can be extended to the case of fractional demands, by solving the appropriate allocation problem within each hexagon. In fact the resulting allocation problem within each hexagon can now be solved *optimally* in polynomial-time using a greedy approach, yielding a truthful auction mechanism with a 7-approximate expected revenue.

Non-zero Minimum Demands. We handle non-zero minimum demands $\{\check{d}_i\}$ by defining a new allocation vector $(y_1, \ldots, y_n)$ wherein y_i is equal to x_i if $x_i \geq \check{d}_i/\hat{d}_i$ and zero otherwise. The arguments of this subsection straightforwardly apply to this new allocation vector.

2) **Service-Provider Based Bidding:** Till now, we have implicitly assumed the base stations (i.e., their demands) are independent. We now consider a more general model considered in [18], wherein base stations belonging to the same service provider bid collectively. More formally, each given base station belongs to a unique service provider, and the demand of each service provider i is given by $(d_{i1}, d_{i2}, \ldots, d_{ij}, \ldots, d_{il}, w_i)$ where d_{ij} is the number of channels required for the j^{th} base station of the i^{th} service provider, l is the total number of base stations for the i^{th} service provider, and w_i is the bid (payment made) if *all* the above demands are satisfied. For simplicity, we assume a unit-disk interference graph between the base stations. Now, to extend our techniques for the above model, we need to assume that the distance between base stations of a service provider is bounded. In other words, all the base stations belonging to a particular service provider i can be enclosed in an R-radius disk centered at a point z_i, where R is a given constant.

Our techniques generalize to the above model as follows. First, as before, we divide the region into hexagons of unit side-length, but use q_2 colors to uniformly color them where

$$q_2 = \min\{x \mid x \geq \max(7, 4R^2/3),\ \text{and}\ x = i^2 + j^2 + ij \\ \text{where}\ i, j \in \mathbb{Z}^+\}. \qquad (8)$$

Using q_2 colors ensures that if a hexagon contains base stations from different service providers i and j, then their corresponding disk-centers z_i and z_j (as defined above) are in hexagons of different colors. The allocation algorithm works as follows.

For each hexagon h, we formulate and solve the following multi-dimensional knapsack (MDKP) problem. Consider the set of hexagons $f(h)$ such that $f(h)$ contains a base station of a service provider i whose disk-center z_i lies in h. The MDKP problem has $|f(h)|$ dimensions, and each dimension has the size-constraint of $M/7$ where M is the total number of available channels. An item of the MDKP is a $f(h)$-dimensional object corresponding to a demand-bid vector of a service provider i whose disk-center z_i lies in h; here, demands for base stations of i belonging to the same hexagon have been aggregated yielding a $f(h)$-dimensional object. Note that

$|f(h)|$ is bounded due to bounded R. Solution to the above MDKP problem yields near-optimal allocation of channels to base stations in $f(h)$ that belong to service provider with disk-centers in h.

We solve the above MDKP problem for each hexagon h in the network region. Then, from the q_2 colors, we pick the color c such that the combination of the MDKP-solutions of the c-colored hexagons yields the most virtual surplus. Below, we prove that the picked allocation is valid and has a $7q_2(1+\epsilon)$-approximate expected revenue.

To show the validity of the returned allocation, we need to prove that: (i) for each hexagon separately, the MDKP-solution is valid, and (ii) the combination of the MDKP-solutions for each color is valid. To show part (i), we divide the M available channels into 7 groups and uniformly distribute them among the hexagons of $f(h)$.[8] Now, we have an MDKP problem where each dimension corresponds to a hexagon of $f(h)$ with the $M/7$ constraint corresponding to the number of channels available in that hexagon.

As for part (ii), it follows from satisfying a modified version of `Property 2` where base stations belonging to two different service providers each falling in a different hexagon with the same color do not interfere. This is satisfied if the distance between any two points in different hexagons of the same color is greater than $2R$. By Lemma 2, this is guaranteed if the number of colors is at least $4R^2/3$. Then, by Lemma 3, the minimum number of colors required would be given by q_2 as defined in Equation 8.

Finally, the $7q_2(1+\epsilon)$-approximate expected revenue follows from similar arguments as in the proof of Theorem 6, except for the fact that we use q_2 colors here (instead of 7) and the extra 7 factor is due to the $M/7$ constraint on each dimension of the MDKP problems.

V. Simulation Results

The main purpose of our simulations is to compare the performance of our mechanisms with other mechanisms in the literature under various settings and performance metrics. Our simulations are split into two parts, one for the TSA-MSW problem and one for the TSA-MER problem. It should be noted here that since no simple auction mechanisms are known for the physical interference model, we restrict our attention to the unit-disk pairwise interference model.

A. Comparing Mechanisms for the TSA-MSW Problem

We start by describing our simulations set-up.

Network Topology and Model. We consider two types of networks, as described below.

- *Random Networks:* We randomly place 50 to 1000 (default being 500) base stations within a fixed area of 1000×1000 square units. The unit-disk interference graph is based on 50-unit radius disks.
- *Real Networks:* We use locations of real cellular base stations available in FCC public GIS database [41] and choose the 843 base stations deployed in the state of Massachusetts. Here, we choose a realistic cell radius of 10 kilometers.

In both networks, we set up an auction of up to 1000 orthogonal single-type channels with the default being 500 channels.

Bidding Functions. We generate general-minded bidding functions for each base station i as follows. First, for each i, we pick l_i (the maximum number of channels bid by i) randomly from the range $[1, m]$, where m is the total number of available channels. Then, we randomly generate i's bid for the first channel and "marginal" bids for each additional channel till l_i. Beyond l_i, marginal bids for each additional channel is assigned zero (to satisfy the free-disposal property). Each marginal bid is chosen from the range $[0, 100]$.

Auction Mechanisms Compared. We compare our auction mechanism for the TSA-MSW problem 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. In addition, we also consider *two* versions of our auction mechanisms: (i) `ours-poly`, based on the polynomial-time allocation algorithm in each hexagon, and (ii) `ours-pseudo-poly`, based on the pseudo-polynomial time algorithm (optimal within each hexagon).

Greedy Auction Mechanism (from [11]). Greedy is a non-truthful mechanism, whose winner determination function allocates channels iteratively to the highest available bid without violating the interference constraint. This allocation results in a 6-approximate social-welfare [11] for the unit-disk model and non-complementary bidding functions, making it the best approximation algorithm known for maximizing social-welfare. If we charge each bidder a payment equal to its bid (declared valuation) for the allocated number of channels, then Greedy's revenue is also within a factor of 6 of the optimal revenue possible.[9] Thus, Greedy's social-welfare is equal to its revenue. Note that Greedy is a pseudo-polynomial algorithm since its running time is polynomial in m, the number of channels, while the other algorithms are polynomial in $\log(m)$.

Naive Auction Mechanism (based on [17]). We now describe a simple auction mechanism (called Naive) that is truthful, but has no performance guarantee on the social-welfare and revenue. Naive's allocation algorithm divides the entire network region into square grid of unit side-length,[10] uniformly colors the resulting square cells using 4 colors, and assigns each color $(1/4)^{th}$ of the available 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

[8]To help visualize this step, imagine that we apply a second layer of coloring on the hexagons of $f(h)$ using 7 colors (see Figure 1). Then, we assign each color one of the 7 groups of $M/7$ channels.

[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.

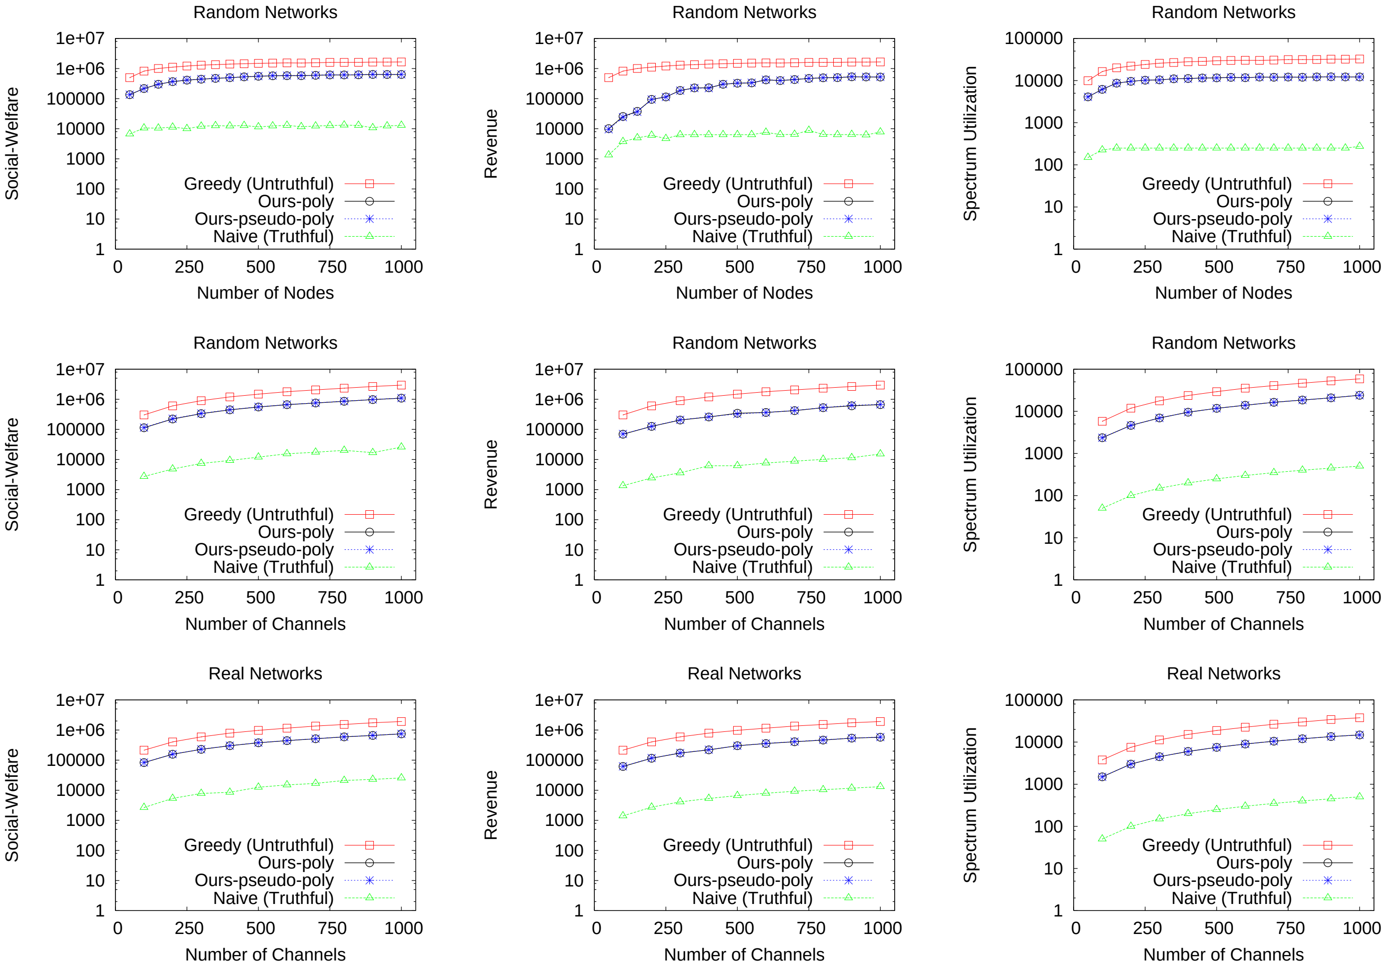

Fig. 3. Performance comparison of various auction mechanisms for the TSA-MSW problem. The first six plots (in the first two rows) are for random networks with varying number of nodes (with 500 channels) and varying number of channels (with 500 nodes). The last three plots are for the cellular network in Massachusetts with 843 base stations and varying number of channels. Recall that `Ours-poly` and `Ours-pseudo-poly` refer to the polynomial-time and the pseudo-polynomial time versions of our auction mechanism. Note that the plots of `Ours-poly` and `Ours-pseudo-poly` overlap with each other due to negligible performance difference. Also, note that the Greedy algorithm runs in pseudo-polynomial algorithm.

H. This is a simple generalization of Vickrey's auction [14] in each square cell.

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. *Spectrum utilization* [17], defined as the total number of allocation pairs in the spectrum allocation, gives a measure of the spatial reuse of a spectrum allocation.

In Figure 3, we plot results for the above three metrics. For the random network, we vary the number of base stations (nodes) as well as the number of available channels, while for the fixed real network we only 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. Finally, we note that the difference in performance of the `ours-poly` and `ours-pseudo-poly` is negligible. This is due to the fact that in practice `ours-poly` performs much better than its worst-case approximation ratio of 1/2 (within a hexagon) for random parameter values, especially when the number of base stations in the hexagon is small and/or the valuations of the bidders are "similar."

Thus, apart from the key properties of truthfulness and provably approximate social-welfare, our auction mechanism also delivers near-optimal revenue.

B. Comparing Mechanisms for the TSA-MER Problem

Network Topology and Model. As with the previous subsection, we consider two types of networks:

- *Random Networks:* We randomly place 100 to 1500 (default being 1000) base stations within a fixed area of 1000×1000 square units. We vary the uniform radius of the disk from 20 to 100 (default being 50) units.
- *Real Networks:* We use locations of real cellular base stations available in FCC public GIS database [41], and choose base stations deployed in 4 different regions:
  - R1: 843 base stations in the state of MA.
  - R2: 2412 base stations in the New England area.

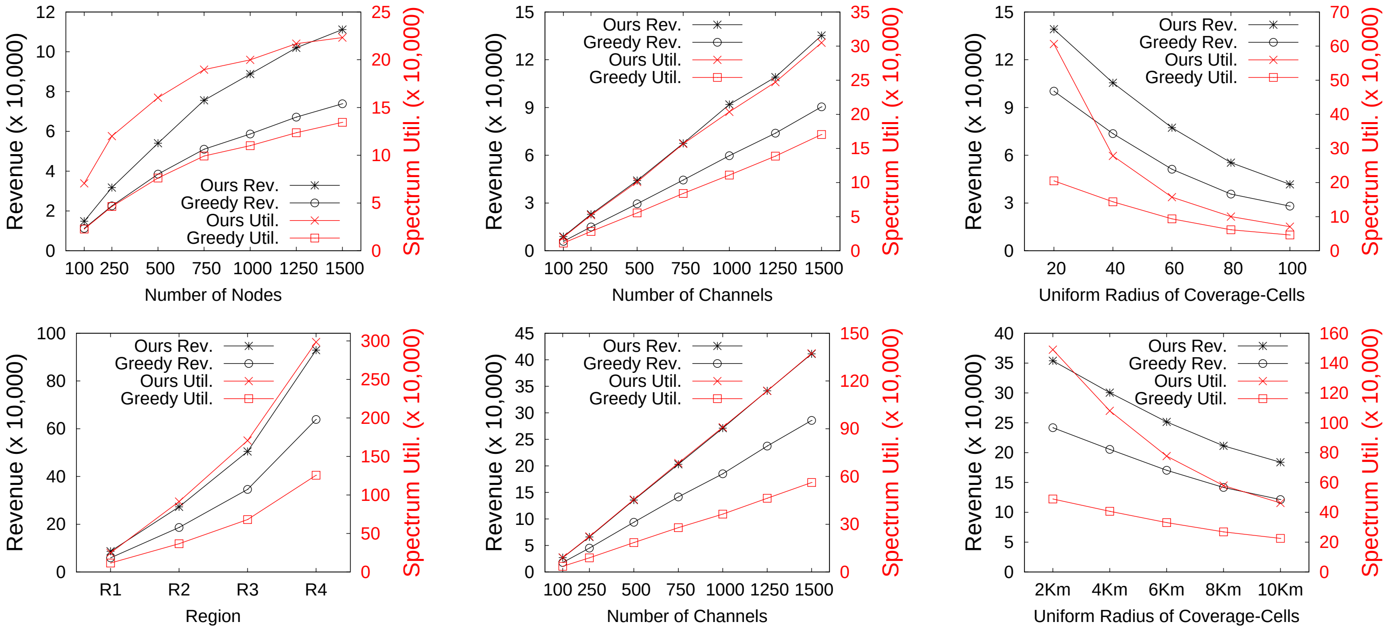

Fig. 4. Performance comparison of our mechanism with the Greedy mechanism for the TSA-MER problem. Generated revenue (left-side y-axis; black solid lines) and the spectrum utilization (right-side y-axis; red dotted lines) on random (top three plots) and real (bottom three plots) networks. The default number of channels is 1000. For random networks, the default number of base stations is 1000, while for real networks, the default region is R2. The default uniform radius of the coverage-cells is 50 units and 5 Km for random and real networks respectively.

– R3: 4467 base stations in New England and NY.
– R4: 8618 base stations in North East USA.

The default region is R2. For all regions, we choose a realistic coverage-cell radius of 5 Kms.

Channels, Demands, and Bids. We set up an auction of up to 1500 orthogonal single-type channels with the default being 1000 channels; this is a reasonable range based on the past FCC auctions [21, 42]. The demands d_i are each chosen randomly from the interval $[1, m]$, where m is the total number of available channels, and the valuations v_i are chosen randomly (and uniformly) from [0,d_i] so that the valuation per channel of each bidder is in the uniform range of [0,1]. For simplicity, we have chosen the valuation-distributions F_i's to be the uniform distributions.

Auction Mechanisms Compared. In our experiments, we compare the below-described enhanced version of our auction mechanism with the Greedy mechanism of [18], the only mechanism in the literature for the TSA-MER problem. The Greedy mechanism is truthful, but has no guarantees on the expected revenue. We note that computing the *optimal* revenue was computationally infeasible even for small networks.

<u>Our Enhanced Mechanism.</u> To further improve the empirical performance of our auction mechanism, we have modified the way we combine the independent solutions of the hexagons. In particular, instead of picking all the hexagons with one of the seven colors, we pick the set of hexagons in a greedy manner as follows. Basically, we pick the hexagons in order of their virtual surplus, while ignoring hexagons that "conflict" with an already picked hexagon; here, two hexagons are considered conflicting if they contain a pair of interfering (and winning) base stations. The above approach certainly yields a valid spectrum allocation. To see the monotonicity (and hence, the truthfulness) of the above approach, note that increasing the bid w_i of a winning bidder i (for fixed bids of other bidders) does not change (a) the presence of i in the FPTAS knapsack-solution, and (b) the winning status of i's hexagon.

Simulation Results. We compare our enhanced auction mechanism with the Greedy [18] mechanism, in terms of the generated revenue and spectrum utilization (total number of allocated channels across all bidders). We conduct experiments for varying: (i) number of base stations, (ii) number of channels, and (iii) the uniform radius of the coverage-cells. See Figure 4. We observe that our mechanism significantly outperforms Greedy in terms of revenue as well as spectrum utilization by an average factor of about 50%, for all parameter values. Moreover, the performance gap generally increases with the increase in the number of channels/base stations or with the decrease in coverage-cells' radius.

Experiments With "Lop-Sided" Demands. In the above experiments with randomly generated demands and bids, our mechanism outperforms the Greedy mechanism by about 50-60%. However, in some cases (as shown in Figure 2), Greedy mechanism can perform arbitrarily bad compared to our mechanism. We now try to generate quasi-random instances, wherein the performance of our mechanism is much better compared to the Greedy mechanism. In particular, we consider randomly generated networks as before, but assign "lop-sided" demands and almost-equal bids to bidders as follows. First, we randomly choose the demands d_i from $([1, \mathcal{I}m] \cup [m - \mathcal{I}m, m])$, where $\mathcal{I}$ is some value between $1/m$ and 1. Then, we assign the low-demand bidders i (i.e., bidders with d_i in $[1, \mathcal{I}m]$) a per-channel bid chosen randomly from [0.95,1]; the (high-demand) bidders get a per-channel bid from [0.9,0.95]. The above assignment of bids is intended to give a slight advantage to the low-demand bidders. Note that, in practice, there is no reason why the bids and demands should have a random distribution. The above specialized setting may reflect a scenario where small start-up concerns compete with

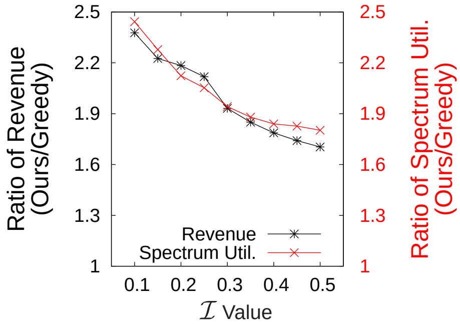

Fig. 5. Performance ratio for "lop-sided" demands in the TSA-MER problem. The demands d_i are randomly chosen from $([1, \mathcal{I}m] \cup [m - \mathcal{I}m, m])$ where $\mathcal{I} \in [0, 1]$. We use random networks of size 1500 base stations with a uniform radius of 50 for the coverage-cells, distributed uniformly in a region of 1000×1000. Number of channels is 1000.

large service providers.

In Figure 5, we show the performance ratio of our mechanism to the Greedy mechanism, in networks of 1500 base stations with coverage-cells of radius 50 units randomly distributed in an region of 1000×1000 units. Number of available channels is 1000. We see that the performance ratio is as high as 2.5, for low values of $\mathcal{I}$, and as expected, the ratio decreases with increase in $\mathcal{I}$.

VI. Conclusions

The recent trend of dynamic spectrum access creates a setting for auctioning pieces of wireless spectrum to competing base stations. To mitigate market manipulation, a truthful spectrum auction is highly desired. We designed truthful spectrum auctions that deliver allocations with either near-optimal expected revenue in the Bayesian setting, or near-optimal social-welfare. We demonstrated the superior performance of our mechanisms through extensive simulations. In future work, we plan to consider more general bidding models.

References

[1] M. Buddhikot et al., "DIMSUMnet: New directions in wireless networking using coordinated dynamic spectrum access," in *WoWMoM*, 2005.

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

[3] V. Brik et al., "DSAP: A protocol for coordinated spectrum access," in *DySPAN*, 2005.

[4] S. Gandhi et al., "A general framework for clearing auction of wireless spectrum," in *DySPAN*, 2007.

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

[6] S. Talukdar et al., "Agents, evolutionary learning and market failure modes," in *HICSS*, 2003.

[7] F. Wolak, "An empirical analysis of the impact of hedge contracts on bidding behavior in a competitive electricity market," *International Economic J.*, 2000.

[8] M. Ilic et al., "Understanding demand: The missing link in efficient electricity markets," Energy Lab. Pub., Tech. Rep., 2001.

[9] C. Barrett et al., "Experimental analysis of bilateral-contracts," *J. of Graph Algos. and Apps.*, 2003.

[10] S. Borenstein et al., "Diagnosing market power in california's restructured wholesale electricity market," NBER, Tech. Rep., 2000.

[11] A. Subramanian et al., "Near-optimal dynamic spectrum allocation in cellular networks," in *DySPAN*, 2008.

[12] O. Ileri et al., "Demand responsive pricing and competitive spectrum allocation via a spectrum server," in *DySPAN*, 2005.

[13] N. Nisan and A. Ronen, "Computationally feasible VCG mechanisms," *J. of AI Research*, vol. 29, 2007.

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

[15] E. Clarke, "Multipart pricing of public goods," *Public Choice*, 1971.

[16] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, 1973.

[17] X. Zhou et al., "eBay in the sky: strategy-proof wireless spectrum auctions," in *MobiCom*, 2008.

[18] J. Jia et al., "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *MobiHoc*, 2009.

[19] Y. Wu et al., "A multi-winner cognitive spectrum auction framework with collusion-resistant mechanisms," in *DySPAN*, 2008.

[20] M. Hoefer et al., "Approximation algorithms for secondary spectrum auctions," in *SPAA*, 2011.

[21] A. Archer et al., "An approximate truthful mechanism for combinatorial auctions with single parameter agents," *Internet Mathematics*, 2003.

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

[23] R. B. Myerson, "Optimal auction design," *Mathematics of Operations Research*, vol. 6, no. 1, 1981.

[24] D. Lehmann, L. I. Oćallaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," *JACM*, 2002.

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

[26] R. Lavi and C. Swamy, "Truthful and near-optimal mechanism design via linear programming," in *FOCS*, 2005.

[27] J. Acharya and R. D. Yates, "Profit maximizing pricing strategies for dynamic spectrum allocation," in *CISS*, 2007.

[28] N. Nisan et al., *Algorithmic Game Theory*. Cambridge Univ. Press, 2007, ch. 13.

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

[30] M. Al-Ayyoub and H. Gupta, "Truthful spectrum auctions with approximate social-welfare or revenue," Stony Brook Univ., Tech. Rep., 2010. [Online]. Available: http://www.cs.sunysb.edu/~hgupta/ps/sw-rev-tech.pdf

[31] ——, "Joint routing, channel assignment, and scheduling for throughput maximization," *TMC*, 2010.

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

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

[34] A. Malakhov and R. Vohra, "Single and multi-dimensional optimal auctions - a network approach," Northwestern Univ., Tech. Rep., 2004.

[35] A. Kumar and G. Iyengar, "Optimal procurement auctions for divisible goods with capacitated suppliers," Columbia Univ., Tech. Rep., 2006.

[36] R. Gautam et al., "Optimal auctions for multi-unit procurement with volume discount bids," in *CEC/EEE*, 2007.

[37] M. Armstrong, "Optimal multi-object auctions," *The Review of Economic Studies*, vol. 67, no. 3, 2000.

[38] J. O. Ledyard, "Optimal combinatoric auctions with single-minded bidders," in *EC*, 2007.

[39] S. Gujar and Y. Narahari, "Optimal multi-unit combinatorial auctions with single minded bidders," in *CEC*, 2009.

[40] H. Kellerer et al., *Knapsack Problems*. Springer, 2004.

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

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

Mahmoud Al-Ayyoub is currently an Assistant Professor in the Computer Science (CS) Department at the Jordan University of Science and Technology. He received a Ph.D. in CS from Stony Brook University, NY, in 2010.

Himanshu Gupta is currently an Associate Professor in the Department of Computer Science at Stony Brook University, New York. His recent research activities focus on theoretical issues in wireless networking, sensor networks, and database systems. He received a B.Tech. in Computer Science and Engineering from IIT, Bombay in 1992, and an M.S. and Ph.D. in Computer Science from Stanford University in 1999.