Truthful Spectrum Auctions With Approximate Revenue

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 in designing an auction mechanism are "truthfulness" and revenue maximization. In this article, we design a polynomial-time spectrum auction mechanism that is truthful *and* yields an allocation with $O(1)$-approximate expected revenue, in the Bayesian setting. Our mechanism generalizes to general interference models. To the best of our knowledge, ours is the first work to design a polynomial-time truthful spectrum auction mechanism with a performance guarantee on the expected revenue. 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 CA 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 *revenue* (total payments by the buyers) [4, 5, 11, 12]. However, 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 lowered revenue. 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 a spectrum auction mechanism that not only encourages truthful behavior but also provides some form of guarantee on the revenue.

Model and Contribution. 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 a polynomial-time auction mechanism that (i) encourages buyers to be truthful (i.e., ensures that the buyers "benefit" the most when their bid is equal to their actual valuation), and (ii) maximizes the generated revenue (sum of the payments by the bidders).

In traditional auction settings, the bidder valuations are completely private. However, when the valuations are private, no truthful auction mechanism can give any performance guarantee on the revenue (see Section II-A). Thus, we consider the relaxed *Bayesian* setting wherein the bidder valuations are drawn from publicly-known probability distributions.

<u>Our Contribution.</u> For the Bayesian setting, we design a polynomial-time spectrum auction mechanism that is truthful and yields $O(1)$-approximate expected revenue. Our mechanism extends to general interference models and other generalizations. To the best of our knowledge, ours is the first work to design a polynomial-time truthful spectrum auction that offers a performance guarantee on the expected revenue.

II. Background and Related Works

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

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; let $\mathbf{w}_i$ denote the *declared valuation function* of the i 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. □

Definition 2: (**Revenue**) 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.

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 $\mathbf{w}_i$. □

Truthfulness and Revenue Maximization. In an untruthful auction, bidders may bid much lower (than their actual valuations) which may indirectly lead to lowered revenue. Thus, unless truthfulness is enforced, maximizing revenue may not be effective. On the other hand, if truthfulness is enforced, then it is not possible [14, 15] to give any guarantees on the generated revenue relative to the optimal revenue, for mechanisms with private valuations. Basically, there is no way to deal with an astronomical bidder even in the simplest case of a single-item auction. Thus, we consider a relaxed setting, known as the *Bayesian* setting, wherein the valuations are drawn independently from publicly-known distributions. Under such a setting, it is possible to design truthful mechanisms with maximum *expected* revenue for simple bidders [16], as described below. In this paper, we essentially extend this classical result to spectrum auctions.

B. Bayesian Setting and Myerson's Optimal Mechanism

In this subsection, we describe the classical Myerson's optimal mechanism for single-parameter auctions in the Bayesian setting. We start with basic definitions.

Single-parameter Auctions. In a *single-parameter* auction, each bidder i has a publicly-known set of outcomes $O_i \subset O$ known as its *winning alternatives* and a private *valuation-value* v_i 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.

Bayesian Setting. In a traditional auction setting, the bidder's valuation is privately-known information which makes it impossible for *truthful* auctions to make any guarantees on the generated revenue [14, 15]. To circumvent this, researchers have considered the *Bayesian setting* wherein each bidder's valuation-value v_i is drawn from a known probability distribution F_i [17].

Myerson's Optimal Mechanism. In a seminal work [16], Myerson presents a truthful optimal mechanism for a single-item auction under the Bayesian setting. Here, we briefly present the key points [17] of Myerson's mechanism applied to the more general single-parameter auctions.

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 [17]. Myerson's mechanism is based on the following characterization of truthful mechanisms for single-parameter auctions.

Theorem 1 ([17, 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 monotone 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 [16], 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 [17] starts by replacing each bid w_i with a *virtual-bid* $\phi_i(w_i)$ as follows.

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

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 2 ([17, 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 1. 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 [17].

Myerson's technique can be easily extended to more general single-parameter auctions [18–20]. Some other works have also extended Myerson's technique to simple multi-parameter settings [21–23].

Applying Myerson's Mechanism To Spectrum Auctions. In a recent work, Jia et al. [24] present a simple extension of Myerson's mechanism 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, [24] presents a polynomial-time mechanism based on the greedy mechanism of [25]. However, the expected revenue delivered by such a mechanism can be arbitrarily bad, as shown in Section III.

Our Work. In this article, we present a polynomial-time truthful spectrum auction mechanism whose expected revenue is within a constant factor of the optimal expected revenue. Our mechanism is based on the above described Myerson's technique, and involves computing an allocation with approximate virtual surplus in polynomial-time.

C. Related Work

Traditional auction mechanism are not directly applicable to spectrum auctions due to the "multi-winner" property of each item (a direct consequence of the spatial reuse of spectrum channels) and wireless interference constraints. Moreover, the corresponding optimization problem of maximizing expected revenue in the context of spectrum auctions is NP-hard (see Section III). Below, we start with discussing recent works on truthful spectrum auctions.

Truthful Spectrum Auctions. To the best of our knowledge, there has been only three works till date, viz., [24, 26, 27], that have designed truthful mechanisms for spectrum auction. We have already discussed [24] in the previous paragraph; we discuss the other two works below.

The truthful mechanism designed by Zhou et al. [26] does not attempt to maximize the revenue or social-welfare. Moreover, their approach is limited to only simple pairwise interference model. As observed in [14, 26], it is rather straightforward to design a truthful auction mechanism without any regard for revenue or social-welfare. However, the authors in [26] 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 [28].

In another work, Wu et al. [27] 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 [29]. Finally, they assume either a single-channel system or that each bidder is interested in only one channel in a multi-channel system.

Other Works. Recently there have been lots of works on dynamic spectrum allocation [4, 5, 11, 12, 30] using either auction-based or pricing-based mechanisms, but all of these works have ignored the truthfulness property.

III. Truthful Spectrum Auction with Approximate Expected Revenue

In this section, we define and address the problem of designing truthful spectrum auctions with approximate expected revenue. 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. Each base station (bidder) i has a publicly-known *demand* for d_i number of channels. The demand d_i determines the winning alternatives for i to be the outcomes wherein i gets at least d_i channels. Each bidder i declares its valuation-value (bid) w_i for the winning alternatives, which may be different than its private valuation-value v_i. We consider the Bayesian setting, wherein the valuation-value v_i is drawn randomly from a publicly-known probability distribution F_i.

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 4: (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.* ***Later, we extend our techniques to pseudo-disk graphs and physical interference model.***

Valid Spectrum Allocation. Given an interference graph and the demands d_i, 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 5: (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 binary allocation vector $(x_1, \ldots, x_{|V|})$ is considered *valid* if there is an assignment $a : V \mapsto P(C)$ such that (i) $|a(i)| \geq d_i$ for all i where $x_i = 1$, 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* to the allocation vector, as is done by the mechanisms designed in this article.

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.

Thus, the TSA-MER problem involves determining (i) a *valid* spectrum allocation, and (ii) payments by each bidder, so that the overall mechanism is truthful and the expected revenue is optimal. 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).

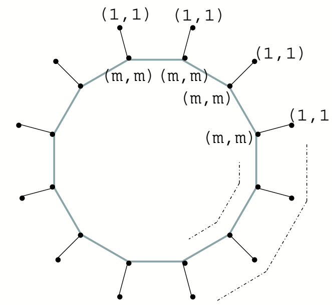

Fig. 1. Counter example for the Greedy mechanism. The figure shows the interference graph over given base stations. The (demand, bid) pair for the “inner” bidders is (m, m), while for the “outer” bidders is $(1, 1)$; here m is the total number of channels. The bids are constant, and hence, virtual-bid 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 $m^2/4$.

Recent Work on TSA-MER. In a recent work, Jia et al. [24] extended Myerson’s mechanism for the TSA-MER problem.[1] However, since maximization of virtual surplus is NP-hard, due to the interference constraint, Myerson’s technique only yields an exponential-time mechanism. Thus, [24] designed a Greedy-heuristic mechanism for the TSA-MER problem as follows. First, the Greedy algorithm sorts the bidders in decreasing order of their virtual-bid per channel (i.e., $\phi_i(w_i)/d_i$). Then, the algorithm considers each bidder in the sorted order, and adds it to the allocation vector if the interference constraint is not violated. Note that to check the violation of 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 1. By Theorem 1, it is easy to show that such a mechanism is truthful. However, the revenue yielded by such a Greedy mechanism can be arbitrarily bad. (see Figure 1).

Below, we design a polynomial-time mechanism that is truthful and yields a valid spectrum allocation with *approximate* expected revenue.

Outline of the Truthful Mechanism with Approximate Expected Revenue. Based on Theorems 1 and 2 of Subsection II-B, 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 1.
2) Determine payments using condition (ii) of Theorem 1.

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 as follows. Basically, we divide the entire network into small hexagonal regions, solve the simpler optimization problem in each hexagon independently, and

[1]We note that [24] 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-B.

then, "combine" the solutions. At a high-level, our algorithm consists of the follows steps.

1) Replace each bid w_i with a *virtual-bid* $\phi_i(w_i)$ as defined by Equation (1).
2) Divide the entire network region into small hexagons of unit side-length.
3) Uniformly-color the hexagons with seven colors.
4) Allocate channels to base stations in each hexagon *independently*, treating it as a Knapsack problem 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) [31] can be used to get a $(1+\epsilon)$-approximate virtual surplus of each hexagon for any $\epsilon > 0$. Note that the interference subgraph in each hexagon is actually a complete graph.
5) For each color, combine the results from all hexagons of that color.
6) Pick the color that has the highest total virtual surplus and allocate the channels to the winners accordingly.
7) Perform a post-processing step to greedily satisfy the demands of more base stations.

The resulting allocation is guaranteed to have at least a $1/7(1+\epsilon)$-factor of the optimal virtual surplus for any $\epsilon > 0$. Moreover, its running time is polynomial in $1/\epsilon$ and the size of the input, i.e., in n and $\log m$, where n is the number of base stations and m is the number of channels.

Plane Division and Coloring. A basic idea in our algorithm is to divide the plane into hexagons of unit side-length (thus, 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.

`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).

Allocation in Each Hexagon. The above properties imply that the channels cannot be re-used inside the same hexagon, but can be re-used across different hexagons of the same color. Thus, allocation in each hexagon can be treated as a Knapsack problem where the virtual-bids are the "values" of items to be placed in the knapsack and the demands are their "weights". The FPTAS of [31] can thus be used to get a $(1+\epsilon)$-approximate virtual surplus of each hexagon for any $\epsilon > 0$.

Combining The Results. Since base stations in different hexagons of same color do not interfere with each other (`Property 2`), 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 virtual surplus. If

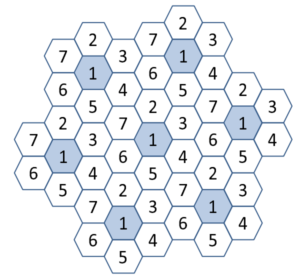

Fig. 2. Hexagons uniformly-colored using 7 colors.

needed, the derived allocation can be easily converted into a channels-to-bidders assignment function.

Post-Processing Step. We will show in Theorem 3 that the above allocation algorithm satisfies the monotonicity of x_i's (i.e., the first condition of Theorem 1). Incidentally, we can further improve the above allocation algorithm, *without violating the monotonicity of x_i's* (as will be shown in Theorem 3), 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 1 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 [24]. 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[2] 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 3: 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 $\epsilon > 0$.

Proof: Truthfulness. By Theorem 1, we need to only show

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

that our allocation algorithm results in monotonically non-decreasing 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 wouldn't 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 doesn't violate the interference constraints, the spectrum allocation returned by the designed mechanism is valid.

Approximate Expected Revenue. By Theorem 2, we only need to show that the virtual surplus of the allocation delivered by the mechanism is within a $7(1+\epsilon)$-factor of the optimal virtual surplus. Since the post-processing step only improves the total virtual surplus, we show below that the allocation before the post-processing step has a $7(1+\epsilon)$-approximate virtual surplus.

Consider a particular color c. For the set of all c-colored hexagons, let I_c be the allocation constructed by our algorithm and O_c be the allocation with the maximum virtual surplus. We show that the virtual surplus of I_c is within a factor of $(1+\epsilon)$ of that of O_c. Note that for any particular hexagon cell, our algorithm constructs an allocation whose virtual surplus is within a factor of $(1+\epsilon)$ of the optimal for that hexagon. Now, by `Property 2`, the virtual surplus of I_c (O_c) is equal to the sum of the virtual surpluses of the constructed (optimal) allocations for the individual c-colored hexagons. Thus, the virtual surplus of I_c is within $(1+\epsilon)$-factor of that of O_c. Finally, the optimal virtual surplus of the full network is upper-bounded by the sum of optimal virtual surplus of each of the seven colors. Since the designed mechanism picks the color with the highest virtual surplus, the virtual surplus of the overall allocation of the designed mechanism is within a factor of $7(1+\epsilon)$ of the optimal. ■

Pseudo-Disk Interference Graphs. We now consider the more general *pseudo-disk interference model*, wherein the coverage cells of base stations may have irregular shapes but are contained within a disk of radius r_1 while containing a disk of radius $r_2 < r_1$. For clarity, we assume that r_1 and r_2 are uniform for all base stations, and r_2 is unit; our techniques can be easily generalized to the non-uniform case using ideas from [32]. For the above pseudo-disk model, our technique works unchanged except that we use q (instead of 7) colors to uniformly[3] color the hexagons, where

$$q = \min\{x | x \geq 4r_1^2/3 \text{ and } x = i^2+j^2+ij \text{ where } i,j \in \mathbb{Z}^+\}.$$

It can be shown [33] that such a q-coloring of hexagons ensures `Property 2`. Thus, the modified mechanism is truthful and yields a $q(1+\epsilon)$ approximate revenue.

IV. Extensions

In this section, we generalize our technique to physical interference model and the service-provider based bidding model. We refer the reader to [33] for other generalization, viz., multi-type channels and fraction-minded bidding model.

A. The Physical Interference Model

We now extend our techniques to the physical interference model. 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 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,

$$\frac{P/\delta_i^\alpha}{\mathcal{N} + \sum_{j \in B'} P/\delta_j^\alpha} \geq \beta, \tag{2}$$

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 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.

Valid Spectrum Allocation. Let $V, C, P(C)$ be the set of bidders, channels, and the power set of channels respectively. 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) $|a(i)| \geq d_i$ for all i where $x_i = 1$, 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., $(P/\delta_i^\alpha)/(\mathcal{N}+\sum_{j \in B} P/\delta_j^\alpha) \geq \beta$

[3]Informally, in a uniform coloring of hexagons, the distance between the "closest" co-colored hexagons is uniform.

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.

Allocation Algorithm. The allocation algorithm for the physical interference model is similar to that for the pairwise interference model, 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.

`Property 3` Every pair of base stations in the same hexagon cannot be concurrently active on the same channel.

`Property 4` 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.

Plane Division and Coloring. It is not hard to see from the SINR Equation (2) that dividing the network region into hexagons of side-length

$$r' = \frac{(\sqrt[\alpha]{\beta}+1)r}{2} \tag{3}$$

would ensure the satisfaction of `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 [32]. Now, to determine appropriate coloring needed to satisfy `Property 4`, we state the following three lemmas; we omit the proof of the first lemma, while the latter two are derived from [29, 34].

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

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

■

Lemma 2: In a hexagonal division with side-length r' and uniformly-colored with x colors, the distance between the centers of a pair 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 4: 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[4] by:

$$q_1 = \min\{x \mid x \geq q_1', x \geq 7, \text{ and } x = i^2+j^2+ij \text{ where } i,j \in \mathbb{Z}^+\} \tag{5}$$

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

■

Overall Allocation Algorithm. As discussed above, dividing the region into hexagons of side-length r' (Equation (3)) and coloring them uniformly using q_1 (Equation (5)) colors, allows us to satisfy `Property 3` and `Property 4`. `Property 3` ensures that allocation in each hexagon can be treated as a Knapsack problem, while `Property 4` allows us to re-use channels across hexagons with the same color. Thus, we can use the same allocation algorithm as in the previous section for the pairwise unit-disk interference, with the above hexagonal division and coloring. Thus, we have the following.

Theorem 5: For the TSA-MER problem under the Bayesian setting and the physical interference model, the above described mechanism is truthful and returns a valid spectrum allocation whose expected revenue is at least $1/(q_1(1+\epsilon))$ of the optimal expected revenue, for a given $\epsilon > 0$. Here, q_1 is as defined in Equation (5). ■

B. 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 [24], wherein base stations belonging to the same service provider bid collectively. More formally, each given base station belongs 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 a 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

$$q_2 = \min\{x | x \geq 4R^2/3 \text{ and } x = i^2+j^2+ij \text{ where } i,j \in \mathbb{Z}^+\}.$$

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

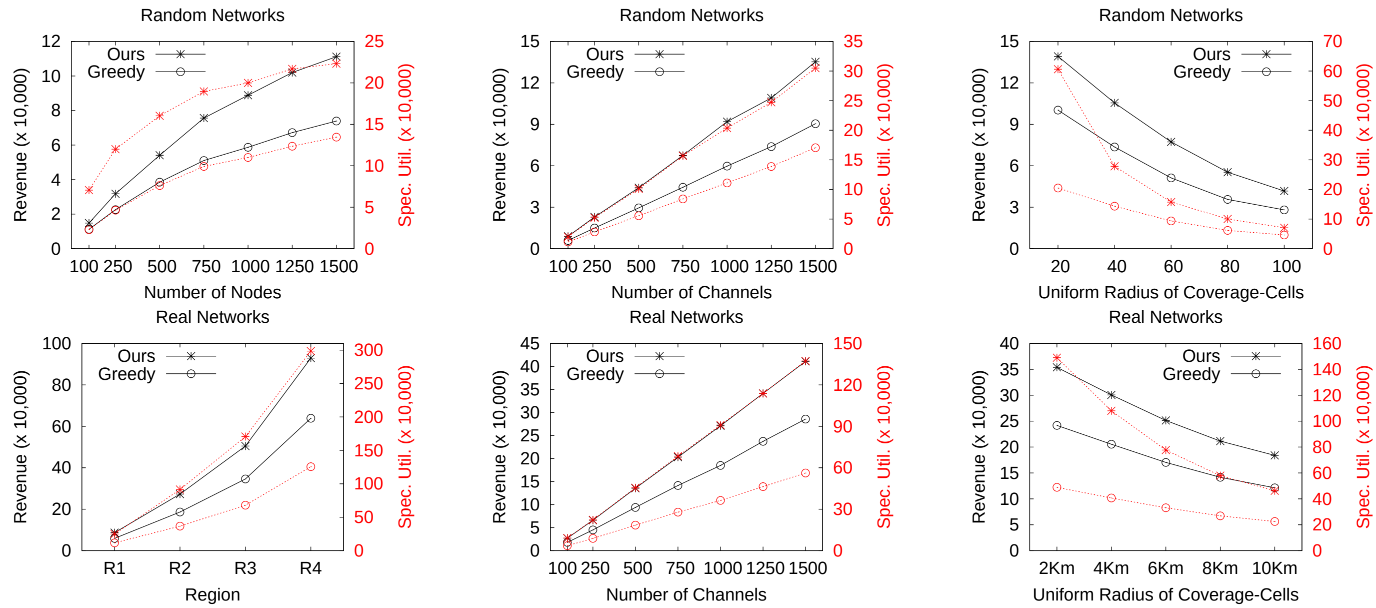

Fig. 3. Generated revenue (left-side y-axis; black solid lines) and the spectrum utilization (right-side y-axis; red dotted lines) on random (first row) and real (second row) networks. The default number of channels is 1000. For random networks, the default number of base stations is 1000, while for 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.

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. It can be shown [33] that the picked allocation is valid (partly, due to the choice of $M/7$ constraint on each dimension of MDKP problems) and has a $7q_2(1+\epsilon)$-approximate expected revenue.

V. Simulation Results

The main purpose of our simulations is to compare the performance of our designed auction mechanism with the Greedy mechanism of [24] under various settings and performance metrics. We start by describing our simulations set-up.

Network. We consider two types of networks:

- *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 100 to 1500 (default being 1000). To generate the interference graph, we use coverage-cells of uniform radius, which is varied from 20 to 100 (default being 50) units.
- *Real Networks:* We use locations of real cellular base stations available in FCC public GIS database [35] and choose base stations deployed in 4 different regions of increasing size and number of base stations.
  - R1: 843 base stations in the state of MA.
  - R2: 2412 base stations in the New England area.
  - 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 [14, 36]. 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 our auction mechanism with the Greedy mechanism of [24], 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.

Simulation Results. We compare our enhanced auction mechanism with the Greedy [24] mechanism, in terms of the generated revenue and spectrum utilization. Here, the *spectrum utilization* is defined as $\sum_i d_i x_i$, the total number of channels allocated across all bidders; it is a measure of the spatial reuse of the spectrum. We conduct experiments for varying: (i) number of base stations, (ii) number of channels, and (iii) the uniform radius of the coverage-cells. We plot our results in Figure 3. 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 the increase in the number of channels/base stations or with the decrease in coverage-cells' radius.

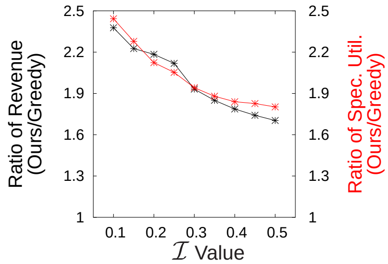

Fig. 4. Performance ratio for "lop-sided" demands. The demands d_i are randomly chosen from $([1, \mathcal{I}m] \cup [m - \mathcal{I}m, m])$ where $\mathcal{I} \in [0, 1]$. We use random networks of size 1500 base stations with a uniform radius of 50 for the coverage-cells, distributed uniformly in a region of 1000×1000. Number of channels is 1000.

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 1), 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, for allocation by the Greedy mechanism. 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 large service providers.

In Figure 4, 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, so that bidders can simply bid their true valuations. In this article, we designed a truthful spectrum auction that delivers an allocation with near-optimal expected revenue, in the Bayesian setting. We have shown the superiority of our mechanism over the Greedy mechanism (the only known mechanism in the literature for our problem) using both theoretical and empirical analysis. In future work, we would 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, M. Chatterjee, and S. Ganguly, "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, Working Papers, 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 Artificial Intelligence Research*, vol. 29, 2007.

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

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

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

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

[18] A. Malakhov and R. Vohra, "Single and multi-dimensional optimal auctions - a network approach," Northwestern University, Discussion Papers, 2004.

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

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

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

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

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

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

[25] D. Lehmann et al., "Truth revelation in approximately efficient combinatorial auctions," *JACM*, 2002.

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

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

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

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

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

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

[32] M. Al-Ayyoub and H. Gupta, "Joint routing, channel assignment, and scheduling for throughput maximization," *IEEE TMC*, 2010.

[33] ——, "Truthful spectrum auctions with approximate revenue," Stony Brook University, Tech. Rep., 2010. [Online]. Available: http://www.cs.sunysb.edu/~hgupta/ps/rev.pdf

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

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

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