Differentially Private and Strategy-Proof Spectrum Auction with Approximate Revenue Maximization

Ruihao Zhu and Kang G. Shin

Department of Electrical Engineering and Computer Science
The University of Michigan, Ann Arbor, MI 48109-2121, U.S.A.
{rhzhu,kgshin}@umich.edu

***Abstract*—The rapid growth of wireless mobile users and applications has led to high demand of spectrum. Auction is a powerful tool to improve the utilization of spectrum resource, and many auction mechanisms have been proposed thus far. However, none of them has considered both the privacy of bidders and the revenue gain of the auctioneer together. In this paper, we study the design of privacy-preserving auction mechanisms. We first propose a differentially private auction mechanism which can achieve strategy-proofness and a near optimal expected revenue based on the concept of virtual valuation. Assuming the knowledge of the bidders' valuation distributions, the near optimal differentially private and strategy-proof auction mechanism uses the generalized Vickrey-Clarke-Groves auction payment scheme to achieve high revenue with a high probability. To tackle its high computational complexity, we also propose an approximate differentially PrivAte, Strategy-proof, and polynomially tractable Spectrum (PASS) auction mechanism that can achieve a sub-optimal revenue. PASS uses a monotone allocation algorithm and the critical payment scheme to achieve strategy-proofness. We also evaluate PASS extensively via simulation, showing that it can generate more revenue than existing mechanisms in the spectrum auction markets.**

I. Introduction

Radio spectrum has become a scarce resource due to the rapid increase in wireless service demand. Conventionally, radio spectrum is allocated in a centralized and static way, but such a policy leaves a large portion of radio spectrum unused in some geographic areas, while making the idle spectrum inaccessible to new wireless application providers that do not have licensed spectrum bands. Consequently, dynamic spectrum allocation (DSA) is introduced to solve or alleviate the problem of spectrum shortage.

Auction is a widely accepted way to tackle the problem of spectrum re-allocation. The Federal Communications Commission (FCC) has adopted it over two decades ago [10]. In recent years, numerous small-scale spectrum auctions have been held, and many spectrum auction mechanisms have also been proposed [36], [37].

There exist three major difficulties in this new spectrum re-allocation. The first difficulty comes from the spatial reusability of each radio spectrum channel. The spatial reusability of a channel can allow two bidders to use the same spectrum simultaneously as long as they are geographically far enough from each other (outside of the interference range). The second is the seller's incentive, which is widely studied for spectrum auction mechanism design [1], [17]. If the seller does not have enough incentive, he will not release his idle channels. The last is strategy-proofness that comprises truthfulness and individual rationality. Intuitively, truthfulness means that no bidder can get a higher payoff by bidding a value other than his true valuation for the goods, while individual rationality means that each bidder gets non-negative utility when bidding truthfully. In the radio spectrum auction market, the bidders are rational, and may manipulate the auction by misreporting to gain benefits. This misreporting, however, may lower the seller's revenue as well as the other bidders' incentive.

Recently, protection of the bid-privacy in spectrum auction has also become an important issue [14], [38]. In most existing studies, a strategy-proof spectrum auction mechanism requires each bidder to report his true valuation, but once the true valuation of a bidder is reported in public, the other bidders can infer the true type of that bidder merely based on the outcome of the auction. In a spectrum auction, a channel is licensed to the bidders for a certain period of time, and all of the bidders should compete again for the usage of the same channel at the end of this period. This makes the inference of a bidder's true type even easier. Moreover, the true type is an important commercial secret for bidders because it can reflect the potential value of the wireless service carried by the spectrum.

When the above four difficulties are taken into account, the problem of designing a privacy-preserving and strategy-proof spectrum auction mechanism for revenue maximization can be very challenging. Previously, cryptography was the main tool for designing privacy-preserving mechanisms [15], [24], but it often incurs high computation and communication overheads, and the performance of the resulting mechanisms may suffer greatly. The recently proposed exponential mechanism, which incorporates techniques originated from differential privacy [23] and mechanism design, provides us a preliminary help to solve the problem. Intuitively, differential privacy means that a single change in the input data set only has limited impact on the output. Therefore, hardly can one make an accurate inference on bidders' bids based on the winners when the exponential mechanism is applied as a large deviation of a single bid does not influence the set of winners much, thereby protecting the bid privacy. To achieve good performance in terms of revenue maximization, however, the exponential mechanism often requires an enumeration of all possible auction outcomes [16], but it is well-

known that even computing the optimal solution for spectrum allocation is NP-complete [6] in multi-hop wireless networks. Besides, exponential mechanisms are often implemented in an approximate truthful manner [23], [38]. Therefore, it is necessary to design proper differentially private mechanisms that can achieve strategy-proofness and approximate revenue maximization.

In this paper, we study the problem of designing a differentially private and strategy-proof auction mechanism for spectrum reallocation. For simplicity of presentation, we consider the case of bidding for a single channel. Each bidder is interested in purchasing a short-term license for this channel in a fixed geographical area. The bidders do not want other bidders and external agents to know their bidding information. Assuming that the seller has a *priori* knowledge of the bidders' valuation distributions, we first design a differentially private and strategy-proof auction mechanism which achieves near optimal expected revenue. This mechanism is based on the exponential mechanism [16]. The near optimal privacy-preserving mechanism can be viewed as a generalization of the Vickrey-Clarke-Groves (VCG) mechanism. Here we adopt the concept of virtual valuation, which is the surplus of the true valuation and a function of valuation distribution, instead of the bidders' original value. This is a widely-used method for the design of conventional mechanisms [13], [18]. It can be shown that maximizing expected revenue is equivalent to maximizing virtual social welfare. The near optimal privacy-preserving mechanism assigns a probability of being chosen for each possible outcome to enforce differential privacy. Nevertheless, this approach is NP-Hard. In order to tackle the problem of high computational complexity, we propose an approximate differentially PrivAte, Strategy-proof, and polynomially tractable Spectrum (PASS) auction mechanism. PASS uses the technique of graph partitioning and the concept of virtual channel to address the spatial reusability of the channel, and a monotone algorithm which combines the features of exponential mechanism and greedy heuristic to allocate the channel. The monotone allocation algorithm, together with the payment rule proposed in [2], guarantees truthfulness. We also implement PASS, and evaluate its performance extensively. This paper makes the following main contributions.

- To the best of our knowledge, this is the first to design differentially private, approximate revenue maximization and strategy-proof mechanism for spectrum auction.
- We model the problem of spectrum reallocation as a sealed-bid auction, and design a near optimal privacy-preserving mechanism. This mechanism is proven to be privacy-preserving, strategy-proof, and achieve a near optimal expected revenue.
- By adopting the graph-partitioning technique [1], we introduce the concept of virtual channel, and propose PASS to reduce the computational complexity of the exponential mechanism. PASS is proven to be a privacy-preserving and strategy-proof auction mechanism, and can achieve a sub-optimal revenue. The computational complexity of PASS is $O(n^2)$, where n is the number of bidders. This low computational complexity makes PASS attractive for short-term lease and large-scale spectrum auctions.
- We implement PASS, and extensively evaluate its performance. Our evaluation results show that PASS achieves differential privacy and a higher revenue than existing approaches.

The remainder of this paper is organized as follows. In Section II, we briefly review the related work in the areas of incentive mechanism design, differentially private mechanisms, and privacy-preserving spectrum auction. Section III presents the model and reviews some related solution concepts, while Section IV details the design of a near optimal privacy-preserving mechanism, and proves its properties. In Section V, we detail the design of PASS and analyze its properties. Section VI presents our evaluation results for PASS. Finally, we conclude the paper with Section VII.

II. Related Work

Numerous efforts have been made to design incentive mechanisms [2], [25], [33]. Nisan [26] studied the general combinatorial auction, while Goldberg *et al.* [11] studied how to auction multiple digital copies. The authors of [19] showed how to construct a mechanism with an approximation algorithm in certain cases via linear programming. The authors of [8] investigated the power of sampling in designing a combinatorial auction mechanism. Dobzinski and Nisan [7] proposed a new method for computing the lower bound for a combinatorial auction mechanism with sub-modular valuations.

McSherry and Talwar [23] first incorporated the techniques of differential privacy into mechanism design and proposed the first differentially private auction mechanism. They used a differentially private pricing approach, but it could only be applied to very simple settings. The algorithmic mechanism design community then focused on how to design truthful and computationally efficient, differentially private auction mechanisms. Leung and Lui [21] studied differentially private and approximately truthful mechanisms in Bayesian setting. Nissim *et al.* [28] modeled privacy loss with disutility functions, and studied truthful mechanisms and their applicability to electronic polling and digital goods' pricing. Chen *et al.* [3] also proposed a model to measure the privacy loss, and studied truthful mechanisms and their applicability to facility allocation and social choice. Nissim *et al.* [29] studied how to convert approximately truthful mechanisms to truthful ones with some privacy loss. In [3], [16], [35], differentially private and truthful mechanisms were studied.

To the best of our knowledge, there only exist a few privacy preserving spectrum auction mechanisms. Huang *et al.* [15] proposed the first strategy-proof and privacy-preserving spectrum auction mechanism. Pan *et al.* [32] also studied the problem of designing a privacy-preserving spectrum auction mechanism to prevent malicious behavior of the auctioneer. However, neither of them makes any performance guarantee in terms of social welfare or revenue. The authors of [14] proposed the first privacy-preserving spectrum auction mechanism with a performance guarantee in social welfare. Nevertheless,

all of the above auction mechanisms used cryptographic tools, and thus induce high computation and communication overheads. Zhu *et al.* [38] proposed the first differentially private spectrum auction mechanism for revenue maximization, but their mechanism achieves only approximate truthfulness.

III. Problem Formulation

In this section, we present the auction model and some related solution concepts.

A. Model

We model the problem of privacy-preserving spectrum allocation as a sealed-bid auction. There is a "seller", who is the primary user. He has a single idle channel,[1] and wants to sublease his idle channel to n secondary users (*e.g.*, wireless service providers), who do not have license to use radio spectrum. The secondary users are the "bidders", and want to buy the license of the idle channel from the seller to provide services to their customers.

In the auction, the seller is trustworthy, and can act as the auctioneer. The seller has a single channel $\mathcal{C}$ for sale, which has the interference range d. Let $\mathcal{N} = \{1, 2, \ldots, n\}$ be the set of bidders. Each bidder $i \in \mathcal{N}$ has a valuation v_i for $\mathcal{C}$, which is private information that we want to protect. The valuation can be calculated as the revenue to be gained by providing wireless services using $\mathcal{C}$. While the exact valuation v_i is private information, we assume that the distribution of v_i, denoted by F_i, is known to the seller, but the bidders do not have information about each other's valuation distribution [17]. Let $f_i(v) = dF_i(v)/dv$ be the corresponding density function. Such information is acquired/inferred from the past transactions. This is known as the *Bayesian setting* for spectrum auction [1]. In the auction, the bidders determine their own bids. Let $\vec{b} = \{b_1, b_2, \ldots, b_n\}$ denote the bid profile which is based on the bid types. Bidders submit sealed bids to the auctioneer simultaneously. We assume that each bidder's valuation and bid are within a given range $[v_{\min}, v_{\max}]$.

We use a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, called *conflict graph*, to describe geographical information, where each node represents a bidder. Any pair of bidders who are separated by a geographical distance smaller than the interference range of $\mathcal{C}$ are said to have "conflict", and are connected via an edge in $\mathcal{G}$. Any pair of bidders who are connected cannot win the channel simultaneously.

Given the bid vector $\vec{b}$ and the conflict graph $\mathcal{G}$, the auctioneer determines the outcome of the auction, denoted by $\vec{x}(\vec{b}) = \{x_1, x_2, \ldots, x_n\}$, where x_i is an indicator *s.t.*,

$$x_i = \begin{cases} 1, & \text{the channel is allocated to } i; \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

The auctioneer also calculates the payment profile, $\vec{p} = \{p_1, p_2, \ldots, p_n\}$, where p_i is bidder i's payment. The auctioneer's revenue, $\mathcal{REV}$, can be computed as the sum of the bidders' payments: $\mathcal{REV} = \sum_{i=1}^{n} p_i$.

[1]This is for the simplicity of presentation, the solution for multiple channels can be easily generated from our solution.

Bidder i's utility u_i is defined as the difference between his valuation times the corresponding indicator and his payment: $u_i = v_i x_i - p_i$. We assume that the bidders are rational and each bidder's goal is to maximize his own utility.

B. Solution Concepts

Here we review some important and useful concepts in mechanism design and differential privacy.

Definition 1 (Dominant Strategy [31])**.** *Strategy a_i is a player i's dominant strategy if for any $a_i' \neq a_i$ and any strategy profile of the other players a_{-i},*

$$u_i(a_i, \vec{a}_{-i}) \geq u_i(a_i', \vec{a}_{-i}).$$

The concept of *truthfulness* is based on that of dominant strategy. Intuitively, truthfulness means that revealing truthful information is the dominant strategy for every player [22].

Definition 2 (Truthful Mechanism [17])**.** *An auction is truthful if and only if any bidder i's (expected) utility of bidding its true valuation v_i is at least its (expected) utility of bidding any other value b_i,*

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

An immediate theorem which relates the monotone algorithm and truthful auction mechanism design [2] follows as:

Theorem 1. *A mechanism is truthful in expectation if and only if for any agent i and any fixed choice of bids by the other agents $\vec{b}_{-i}$,*

1) *$x_i(\vec{b})$ is monotone nondecreasing in b_i;*
2) *$p_i(\vec{b}) = b_i y_i(\vec{b}) - \int_0^{b_i} y_i(z)dz$, where $y_i(z)$ is the probability that bidder i is selected when his bid is z.*

Let's review concepts related to differential privacy. Informally, differential privacy means that the outcome of two nearly identical input data sets (differring in a single input) should also be nearly identical. Formally,

Definition 3 (Differential Privacy [9])**.** *A randomized computation M has ε-differential privacy if for any two input sets A and B with difference in a single input, and for any set of outcomes $R \subseteq RANGE(M)$,*

$$\mathbf{Pr}[M(A) \in R] \leq \exp(\varepsilon) \times \mathbf{Pr}[M(B) \in R].$$

One important property of differential privacy is composability:

Corollary 1 (Composability [23])**.** *The sequential application of randomized computation M_i, each giving ε_i-differential privacy, yields $\sum_i \varepsilon_i$ differential privacy.*

There is also a relaxed definition of differential privacy:

Definition 4 (Approximate Differential Privacy [9])**.** *A randomized computation M has (ε, δ)-differential privacy if for any two input sets A and B with a single data difference, and for any set of outcomes $R \subseteq RANGE(M)$,*

$$\mathbf{Pr}[M(A) \in R] \leq \exp(\varepsilon) \times \mathbf{Pr}[M(B) \in R] + \delta.$$

Incorporating differential privacy, one powerful tool in mechanism design is the exponential mechanism [16], [23]. Mapping the input data set A and an outcome r in the outcome space R to a certain score function $q(A,r)$, the exponential mechanism $\varepsilon_q^\varepsilon(A)$ satisfies: $\mathbf{Pr}[\varepsilon_q^\varepsilon(A)=r] \propto \exp(\varepsilon q(A,r))$.

This exponential mechanism guarantees a $2\varepsilon\Delta$-differential privacy, where Δ is an upper-bound of difference of two data sets. An immediate theorem can also be derived as [12]:

Theorem 2. *When used to select an output $r \in R$, the exponential mechanism $\varepsilon_q^\varepsilon(A)$ yields $2\varepsilon\Delta$-differential privacy. Let R_{OPT} denote the subset of R achieving $q(A,r) = max_r q(A,r)$, then the exponential mechanism ensures that*

$$\mathbf{Pr}\left[q(A,\varepsilon_q^\varepsilon(A)) < \max_r\ q(A,r) - \frac{\ln(|R|/|R_{OPT}|)}{\varepsilon} - \frac{t}{\varepsilon}\right] \leq \exp(-t). \tag{3}$$

The goal of our auction mechanism design is to achieve strategy-proofness, privacy preservation and revenue maximization. The problem of revenue maximization can be formulated as a binary programming problem as:

Objective:

$$\text{Maximize} \quad \mathcal{REV} = \sum_{i=1}^{n} p_i x_i$$

Subject to:

$$x_i + x_j \leq 1, \qquad \forall (i,j) \in \mathcal{E}; \tag{4}$$

$$x_i \in \{0,1\}, \qquad \forall i \in \mathcal{N}, \tag{5}$$

but this is known to be computationally intractable [34], nor is it privacy-preserving. Therefore, we need to seek an approximation approach.

IV. Near Optimal Privacy-Preserving Mechanism

We now present a near optimal privacy-preserving mechanism. This mechanism is based on the exponential mechanism proposed in [16], which is a strategy-proof mechanism that maximizes the auctioneer's expected free social welfare (see [16] for the meaning of free social welfare). We first adopt the concept of virtual valuation [27], which is essential for revenue maximization. We apply the exponential mechanism on the virtual valuations, which guarantees the maximum expected free revenue while enforcing strategy-proofness and ε-differential privacy.

A. Virtual Valuation and Virtual Surplus

We first adopt the concept of virtual valuation and virtual surplus from [27]:

Definition 5. *The virtual valuation of agent i with valuation v_i is:*

$$\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}. \tag{6}$$

A company concept is the *virtual bid*, which is calculated by plugging in the bid into Eq. (6). The virtual bid profile of the bidders is denoted by $\vec{\phi}(\vec{b}) = \{\phi_1(b_1), \phi_2(b_2) \ldots, \phi_n(b_n)\}$.

Definition 6. *Given valuations, v_i , and the corresponding virtual valuations, $\phi_i(v_i)$, the virtual surplus of allocation $\vec{x}$ is:*

$$\sum_{i=1}^{n} \phi_i(v_i) x_i. \tag{7}$$

We further assume that the distributions of the bidders satisfy the *monotone hazard rate* (*i.e.*, $f_i(v)/(1-F_i(v))$ is monotone non-decreasing), so that the virtual valuations are monotone non-decreasing. This is a sufficient condition for a strategy-proof mechanism [27]. We also assume that for each bidder i, his virtual valuations are bounded by $[\phi_i(v_{\min}), \phi_i(v_{\max})]$, and the difference between $\phi_i(v_{\max})$ and $\phi_i(v_{\min})$ is denoted by Δ.

An immediate theorem is that any truthful mechanism has an expected revenue equal to its expected virtual surplus:

Theorem 3. *The expected profit of any truthful mechanism is equal to its expected virtual surplus:*

$$\mathbf{E}[\mathcal{REV}] = \mathbf{E}\left[\sum_{i=1}^{n} \phi_i(v_i) x_i\right]. \tag{8}$$

Due to space limitation, we omit the proof here, but the proof of this theorem can be found in [27].

B. Design Details

With the help of virtual valuations and virtual surplus, the problem of designing privacy-preserving and near optimal maximization spectrum auction can be solved with the exponential mechanism. Following the guiding principle of [16], the mechanism assigns each outcome a probability, which is proportional to the exponential of its revenue, and then selects an outcome accordingly. The mechanism also charges each bidder a VCG-like payment to enforce strategy-proofness. The detailed auction works as follows.

1) For each bidder i, calculate the virtual bid

$$\phi_i(b_i) = b_i - \frac{1-F_i(b_i)}{f_i(b_i)}. \tag{9}$$

2) Select an outcome $\vec{x}$ satisfies Eqs. (4)and (5) with probability

$$\mathbf{Pr}[\vec{x}] \propto \exp\left(\frac{\varepsilon}{2\Delta}\sum_i \phi_i(b_i)\cdot x_i\right). \tag{10}$$

3) The payment for a winner bidder i is $p_i = \phi_i^{-1}(p_i')$, where

$$\begin{aligned} p_i' = &- \mathop{\mathbf{E}}_{\vec{x}\sim EXP_\varepsilon^R(\phi_i(b_i),\vec{\phi}_{-i}(\vec{b}_{-i}))}\left[\sum_{k\neq i}\phi_k(b_k)x_k\right] \\ &- \frac{2\Delta}{\varepsilon}\cdot S\left(EXP_\varepsilon^R(\phi_i(b_i),\vec{\phi}_{-i}(\vec{b}_{-i}))\right) \\ &+ \frac{2\Delta}{\varepsilon}\cdot \ln\left(\sum_{\vec{x}}\exp\left(\frac{\varepsilon}{2\Delta}\sum_{k\neq i}\phi_k(b_k)x_k\right)\right), \end{aligned} \tag{11}$$

where

$$\vec{b}_{-i} = \{b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n\} \tag{12}$$

$$\vec{\phi}_{-i} = \{\phi_1(b_1), \ldots, \phi_{i-1}(b_{i-1}), \phi_{i+1}(b_{i+1}), \ldots, \phi_n(b_n)\} \tag{13}$$

is the virtual bid profile of all bidders other than i, EXP^R_ε is the assigned probability distribution over the virtual bid profile of all bidders in Step 2, and $S(\cdot)$ is the Shannon entropy [5].

C. Analysis of the Near Optimal Privacy-Preserving Mechanism

Here we analyze the properties of the near optimal privacy-preserving mechanism. We first show that our design can achieve near optimal expected revenue; we then show that together with our VCG-like payment design, the mechanism achieves strategy-proofness; we finally show that the mechanism achieves ε-differential privacy.

1) Near Optimal Expected Revenue

By Theorem 3, we show that the expected revenue of the mechanism is equivalent to the expected virtual surplus. Therefore, we analyze the expected virtual surplus. One important concept that can help us is the "free energy" in the literature of physics and chemistry [30]. Similarly, we first define the *free virtual surplus* ($\mathcal{FVS}$)

$$\mathcal{FVS}(\mathcal{F}) = \mathop{\mathbf{E}}_{\vec{x}\sim\mathcal{F}}\left[\sum_{k=1}^{n}\phi_k(b_k)\cdot x_k\right] + \frac{2\Delta}{\varepsilon}\cdot S(\mathcal{F}), \tag{14}$$

where $\mathcal{F}$ is a distribution over $\vec{x}$. We prove the following theorem, which states that when the distribution $\mathcal{F}$ in Eq. (14) is chosen as in Step 2 of the near optimal mechanism, the free virtual surplus is maximized over the input bid profile of the bidders

Theorem 4. *Given bid vector $\vec{b}$, $\mathcal{FVS}$ is maximized when*

$$\mathcal{F} = EXP^R_\varepsilon(\vec{\phi}(\vec{b})),$$

and the maximum value is

$$\frac{2\Delta}{\varepsilon}\sum_{\vec{x}}\ln\left(\sum_{\vec{x}}\exp\left(\frac{\varepsilon}{2\Delta}\langle\vec{\phi}(\vec{b}),\vec{x}\rangle\right)\right), \tag{15}$$

where $\vec{x}\sim EXP^R_\varepsilon(\phi_i(b_i),\vec{\phi}_{-i}(\vec{b}_{-i}))$.

Due to space limitation, we omit the proof here; see our technical report for the details of the proof [39].

Now that the free virtual surplus is equal to the expected revenue plus a term, the mechanism achieves a near optimal expected revenue.

2) Strategy-proofness

With the above theorem, we are now ready to prove that the near optimal privacy-preserving mechanism achieves truthfulness. Recall that in the VCG auction [4], each winning bidder is charged the externality he exerts on the rest of the bidders (*e.g.*, the payment of a winner i is the difference between the optimal social welfare of the bidders other than i and the optimal social welfare of all the bidders minus i's bid); otherwise, the bidder is charged 0.

We start by proving that the near optimal privacy-preserving mechanism achieves truthfulness.

Lemma 1. *The proposed mechanism achieves truthfulness.*

We then show that the near optimal privacy-preserving mechanism achieves individual rationality.

Lemma 2. *The proposed mechanism achieves individual rationality.*

Due to space limitation, we omit the proof; see our technical report for the details of the proof [39].

3) Differential Privacy

We finally state that the near optimal privacy-preserving mechanism can preserve the bidders' valuation privacy.

Theorem 5. *The near optimal privacy-preserving mechanism preserves ε-differential privacy for bidders' valuation privacy.*

This theorem is a corollary of Theorem 2.

V. PASS: A Differentially Private and Strategy-Proof Spectrum Auction Mechanism

The near optimal auction mechanism introduced in the previous section can protect the bidders' bid privacy as well as generate a near optimal expected revenue. Nevertheless, its computational complexity is prohibitively high, thus making it unsuitable for large-scale spectrum markets. Therefore, by adopting the graph partitioning technique [1], and introducing the concept of virtual channel, we propose PASS, which is a polynomially tractable, differentially private, and strategy-proof spectrum auction mechanism. PASS achieves good performance in terms of approximate revenue maximization.

A. Design Rationale

PASS integrates the exponential mechanism with the greedy heuristic for a single-minded combinatorial auction mechanism [20] into spectrum auction to achieve both approximate revenue maximization and differential privacy. Basically, PASS chooses the winning bidders iteratively. In each iteration, each remaining bidder is assigned a probability of being chosen, and PASS chooses one of them as a winning bidder. The main idea of PASS is to first apply the technique of graph partitioning, and create a set of virtual channels that capture the conflict among the bidders for each bidder, thus transorming the spectrum auction into a single-minded combinatorial auction like scenario. Finally, PASS computes the probability of being chosen for each bidder based on a specific norm, and chooses the set of winning bidders iteratively, and determines the payment for each bidder.

B. Design Details

Following the guidelines in Section V-A, we describe PASS in detail. PASS performs the auction in four steps. It first partitions the conflict graph into uniform hexagons, thus dividing the bidders into groups. Afterwards, PASS creates virtual channels for bidders according to their geographical locations and the conflict graph. PASS then computes the probability of being chosen for each bidder, and chooses winning bidders randomly based on the assigned probability iteratively. Finally, PASS determines the payment for each bidder.

Phase 1: Graph Partition

PASS first partitions the geographical area [1]. Since the interference range of $\mathcal{C}$ is d, PASS divides the entire area into small hexagons with side-length equal to half of the interference range of $\mathcal{C}$ (*i.e.*, the side-length is $d/2$). Suppose there are a total of m hexagons, and correspondingly, m bidder groups, denoted by $\mathcal{A} = \{a_1, a_2, \ldots, a_m\}$. Each bidder i belongs to exactly one hexagon a_j (*i.e.*, $\bigcup_{j=1}^{m} a_j = \mathcal{N}$). An illustrative example of the partition with $m = 24$ is shown Fig. 1.

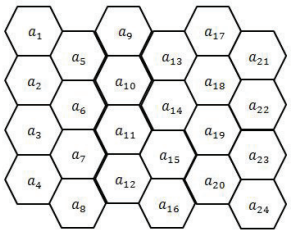

Fig. 1. An illustrative example of partition.

The partition satisfies that the maximum distance within a single hexagon is less than or equal to d (*i.e.*, the diametrical length) so that all the bidders in a single hexagon are mutually conflict (*i.e.*, the conflict graph over a single hexagon is a complete graph).

Fact 1. *Any pair of bidders from the same hexagon are conflicting bidders, and can not be allocated to the channel simultaneously.*

The graph partition technique is the basis of virtual channel, and the intuition behind this partition is that in the single-minded combinatorial auction, the approximation ratio is equal to $\sqrt{k}$ [20], where k is the maximum size of the bundles of all the bidders; while in PASS, the maximum size of the bundles is determined by the number of virtual channels each bidder is interested in. By partitioning the geographical area, the number of virtual channels each bidder is interested in is cut down.

Phase 2: Virtual Channel Assignment

With the partitioned graph, PASS introduces virtual channel to capture the interference among the bidders. Specifically, a virtual channel $vc_{j,k}$ (j and k can be the same) is assigned to bidder i if he satisfies one of the following two conditions:

1) i locates in a_j, and he is in conflict with at least one bidder i' in a_k (*i.e.*, i and i' cannot be granted the channel simultaneously).
2) i locates in a_k, and he is in conflict with at least one bidder i' in a_j (*i.e.*, i and i' cannot be granted the channel simultaneously).

For convenience, let $\vec{r} = \{r_1, r_2, \ldots, r_n\}$ denote the set of virtual channels assigned to each bidder, where r_i is the bundle of virtual channels assigned to bidder i. In other words, i is interested in the bundle of virtual channels r_i. The formal process of assigning virtual channels to each bidder is described in Algorithm 1, where $DISTANCE(i,i')$ is the geographical distance of bidders i and i'. The assignment of virtual channels guarantees that the intersection of the sets of virtual channels assigned to a pair of conflict bidders is non-empty.

Lemma 3. *For any pair of bidders i and j, if i and j are in conflict* (i.e., $(i.j) \in \mathcal{E}$), *then $r_i \cap r_j \neq \emptyset$.*

Algorithm 1 Virtual Channel Assignment

Input: A conflict $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$. A partitioned groups $\mathcal{A}$.
Output: A vector of bundles of virtual channels each bidder is interested in $\vec{r}$.

```
1: r⃗ ← ∅⃗.
2: for all a_j ∈ A, a_k ∈ A do
3:    for all i ∈ a_j, i'(≠ i) ∈ a_k do
4:       if DISTANCE(i, i') ≤ d then
5:          r_i ← r_i ∪ {vc_{j,k}}, r_{i'} ← r_{i'} ∪ {vc_{j,k}}.
6:       end if
7:    end for
8: end for
9: return r⃗.
```

Phase 3: Winner Determination

Based on the partitioned graph and each bidder's interested virtual channel. PASS first calculates each bidder's virtual bid, and then assigns each bidder i a probability of being chosen with respect to a specific norm, which is the bidder's virtual bid over the square root of the size of the set of virtual channels he is interested in, *i.e.*,

$$\frac{\phi_i(b_i)}{\sqrt{|r_i|}}. \tag{16}$$

PASS then chooses the winning bidders iteratively, and maintains a set of remaining bidders, denoted by $\mathcal{R}$. Initially, $\mathcal{R} = \mathcal{N}$. In each iteration, for all the bidders $i \in \mathcal{R}$, the probability of being chosen is proportional to the exponential of the norm defined in Eq. (16) times a constant; otherwise, the probability is 0, *i.e.*,

$$\mathbf{Pr}[\mathcal{W} \leftarrow \mathcal{W} \cup \{i\}] \propto \begin{cases} \exp\left(\varepsilon' \frac{\phi_i(b_i)}{\sqrt{|r_i|}}\right) & \text{if } i \in \mathcal{R}, \\ 0, & \text{otherwise.} \end{cases} \tag{17}$$

where $\varepsilon' = \varepsilon / (e\Delta \ln(e/\delta))$ and $\mathcal{W}$ is the set of winning bidders. PASS then normalizes the values and chooses a winning bidder accordingly, *i.e.*,

$$\mathbf{Pr}[\mathcal{W} \leftarrow \mathcal{W} \cup \{i\}] = \begin{cases} \frac{\exp\left(\varepsilon' \cdot \phi_i(b_i)/\sqrt{|r_i|}\right)}{\sum_{j \in \mathcal{R}} \exp\left(\varepsilon' \cdot \phi_j(b_j)/\sqrt{|r_j|}\right)} & \text{if } i \in \mathcal{R}, \\ 0, & \text{otherwise.} \end{cases} \tag{18}$$

Suppose bidder i is chosen as a winner. Then, PASS removes i from $\mathcal{R}$. If bidder $j \in \mathcal{R}$ is in conflict with i (*i.e.*, $r_i \cap r_j \neq \emptyset$), then j is also removed from $\mathcal{R}$. PASS repeats this until there is no bidder left in $\mathcal{R}$. Selection of the winner is described formally in Algorithm 2.

PASS takes $O(n^2)$ time to assign virtual channels. For the winner determination, PASS takes $O(n)$ time to assign the probability of being chosen to each remaining bidder in an iteration. Since there are at most n iterations, PASS

takes $O(n^2)$ to determine the winners. Therefore, the total computational complexity of PASS is $O(n^2)$.

Phase 4: Payment Scheme

PASS's winner determination algorithm is monotonic (see the next subsection for a proof). According to Theorem 1, the payment for a bidder i is

$$p_i(\vec{b}) = b_i y_i(\vec{b}) - \int_0^{b_i} y_i(z)dz, \tag{19}$$

where $y_i(z)$ is generalized to be the probability that bidder i wins the channel when his bid is z.

Algorithm 2 Differentially Private and Strategy-Proof Spectrum Auction Mechanism

Input: A conflict graph $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$, a virtual bid vector $\vec{\phi}(\vec{b})$, and a vector of virtual channel requests $\vec{r}$

Output: A set of winners $\mathcal{W}$.

$\mathcal{W} \leftarrow \emptyset, \varepsilon' \leftarrow \varepsilon/(\Delta \cdot e\ln(e/\delta)), \mathcal{R} \leftarrow \mathcal{N}$.
while $\mathcal{R} \neq \emptyset$ **do**
  for all $i \in \mathcal{R}$ **do**
    $\mathbf{Pr}[\mathcal{W} \leftarrow \mathcal{W} \cup \{i\}] = \frac{\exp\left(\varepsilon' \cdot \phi_i(b_i)/\sqrt{|r_i|}\right)}{\sum_{j \in \mathcal{R}} \exp\left(\varepsilon' \cdot \phi_j(b_j)/\sqrt{|r_j|}\right)}$.
  end for
  Select i according to the computed probability distribution.
  if i is selected **then**
    $\mathcal{R} \leftarrow \mathcal{R} \setminus \{i\}$.
    for all $j \in \mathcal{R}$ **do**
      if $r_j \wedge r_i \neq \emptyset$ **then**
        $\mathcal{R} \leftarrow \mathcal{R} \setminus \{j\}$.
      end if
    end for
  end if
end while
return $\mathcal{W}$.

C. Analysis of PASS

The properties of PASS are analyzed in this subsection.

1) Revenue

We first analyze the revenue guarantee of PASS. One of our key steps is to provide an upper-bound for the size of the virtual channels bundle assigned to each bidder.

Lemma 4. *The size of the virtual channels bundle assigned to each bidder is less than or equal to 12:*

$$|r_i| \leq 12 \ \forall i \in \mathcal{N}. \tag{20}$$

Proof. Consider the hexagon a_{10} in Fig. 2 as an example. Due to its topological symmetry, we divide a_{10} into 12 identical right triangles. We take one of them and denote it by T (as shown in the shaded area in a_{10}). After some simple calculation, we can conclude that all the bidders in conflict with any of the bidders in T lie in the 12 colored hexagons, and no bidder from other hexagon will be in conflict with bidders in T. An arbitrary bidder i must be in one of the right triangles, and hence $|r_i| \leq 12$. □

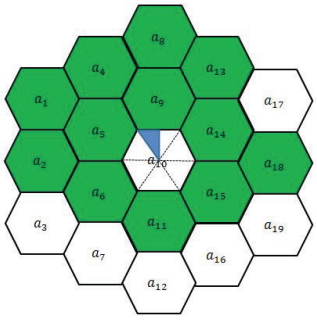

Fig. 2. Bounding r_i.

We are now ready to prove the revenue guarantee of PASS:

Theorem 6. *With the probability of at least $1 - 1/n^{O(1)}$, PASS can generate a set of winners with a revenue of at least $\mathcal{REV}^*/12 - O(\ln n)$. Here $\mathcal{REV}^*$ is the optimal revenue.*

Due to space limitation, we omit the proof; see our technical report for the details of the proof [39].

2) Strategy-proofness

Here we prove that PASS achieves strategy-proofness. According to Theorem 1, what we need to prove is that the allocation rule of PASS is monotone (*i.e.*, for each bidder i, the probability that he is selected as a winner is monotonically non-decreasing with his bid).

Lemma 5. *For each bidder i, the probability that he is selected as a winner is monotonically non-decreasing with his bid b_i.*

Due to space limitation, we omit the proof; see our technical report for the details of the proof [39].

By the above Lemmas and Theorem 1, we can conclude that PASS achieves strategy-proofness.

Theorem 7. *PASS achieves strategy-proofness.*

3) Differential Privacy

We finally prove that PASS can preserve the bidders' valuation privacy.

Theorem 8. *For any $\delta \leq 1/2$, PASS preserves $(\varepsilon'(e-1)\Delta\ln(e/\delta), \delta)$ differential privacy for bidders' virtual bids.*

Due to space limitation, we omit some intermediate steps; see our technical report for the details of the proof [39].

VI. Numerical Results

We have implemented PASS (in a simulator) and extensively evaluated its performance. Our evaluation results show that PASS can not only generate a relatively high revenue, but also achieves good differential privacy.

A. Methodology

To evaluate the performance of PASS in terms of revenue maximization, we compare PASS with DEAR, which is a differentially private and approximately truthful mechanism

for spectrum auction proposed in [38] and a truthful spectrum auction mechanism (denoted as "TSAWAP") proposed in [1].

The number of bidders is varied from 100 to 1500 with a step of 100, and the bidders are randomly deployed in a square area of 1000m$\times$1000m. Each bidder has an interference range of 425m [38]. We assume that each bidder's bid follows a uniform distribution over [0,1]. We vary the number of channels from 5 to 15 with a step of 5. We set the privacy constant ε to 0.1 and 0.5, and δ to 0.25.[2] All the results are averaged over 1000 runs.

We use two metrics to evaluate the performance of PASS—–*revenue* and *privacy*. Revenue refers to the sum of charges to the bidders. A mechanism guarantees good privacy if the probability distribution over an arbitrary outcome has as small a change as possible when any bidder unilaterally reports a different bid. Following the definition of differential privacy, we definite the the notion of *Privacy Leakage* (note that this is different from the one proposed in [38]) to quantitatively measure the privacy guarantee of PASS:

Definition 7 (Privacy Leakage)**.** *Given a mechanism $\mathcal{M}$, let $\vec{\phi}$ and $\vec{\phi}'$ be virtual bid vectors for bidding profiles $\vec{b}$ and $\vec{b}'$, which only differ in a single entry, respectively. Let O be the outcome space The privacy leakage between the two bidding profiles is the maximum of absolute differences between the logarithmic probabilities of any outcome,* i.e.*,*

$$\max_{o\in O} |\ln \pi_o - \ln \pi'_o|, \tag{21}$$

where π and π' is the probability distribution over the outcome space with respect to $\vec{\phi}$ and $\vec{\phi}'$, respectively.

B. Revenue

We first evaluate PASS's performances in terms of revenue.

Fig. 3(a) shows the comparison of PASS with DEAR and TSAWAP, when the number of channels is 5. From these results, PASS is shown to outperform DEAR and TSWAP in nearly all the cases in terms of revenue. The only exception is when the number of bidders is 500 (for both $\varepsilon = 0.1$ and $\varepsilon = 0.5$), DEAR can generate slightly more revenue than PASS. This is because PASS relies on the existence of critical neighbors to generate revenue. When the number of bidders grows, PASS can find critical neighbors with higher bids.

Fig. 3(b) compares PASS with DEAR and TSAWAP, when the number of channels is 10. From these results, PASS is outperforms DEAR and TSWAP in nearly all of the cases in terms of revenue. The only exception is when the number of bidders is less than 1000 (for both $\varepsilon = 0.1$ and $\varepsilon = 0.5$), DEAR can generate more revenues than PASS. This is because when the number of channels increases, PASS cannot find enough critical neighbors to generate revenue. When the number of bidders grows, it is easier for PASS to find critical neighbors.

Fig. 3(c) shows the comparison results of PASS with DEAR and TSAWAP, when the number of channels is 15. From these results, PASS is shown to outperform DEAR and TSWAP in most of the cases in terms of revenue. The only exception are when the number of bidders is less than 1000 (for $\varepsilon = 0.5$), and when the number of bidders is than 1500 (for $\varepsilon = 0.1$), DEAR can generate more revenue than PASS. This is also because PASS relies on the existence of critical neighbors to generate revenue. When the number of channels is large, PASS cannot find enough critical neighbors.

All of the above figures show that PASS outperforms TSAWAP because TSAWAP can only use critical neighbors in the same hexagon to generate revenue while PASS can use critical neighbors in other hexagons to generate revenue. This leads to critical neighbors with higher bids.

In summary, PASS is shown to be suitable for large-scale secondary spectrum markets, especially when spectrum channels are scarce.

C. Privacy

Next, PASS is evaluated in terms of privacy preservation.

Fig. 4(a) shows the privacy leakage of PASS when $\varepsilon = 0.1$ and $\varepsilon = 0.5$ (under this setting, PASS is supposed to achieve $((e-1)/10e, 0.25)$ and $((e-1)/2e, 0.25)$ differential privacy), respectively. The number of channels is 5. The results show that PASS's privacy leakage is always less than 0.04 when $\varepsilon = 0.1$, and it is always less than 0.15 when $\varepsilon = 0.5$. The results also show that when the number of bidders increases, the privacy leakage of PASS decreases. This is because with more bidders, the probability of each outcome becomes smaller. Besides, we can see that the privacy performance of PASS is far better than $((e-1)/10e, 0.25)$ and $((e-1)/10e, 0.25)$ differential privacy when $\varepsilon = 0.1$ and 0.5, respectively. Therefore, it is nearly impossible for any agent to learn the bid information of the bidders when PASS is implemented.

Figs. 4(b) and 4(c) show the privacy leakage of PASS, when the number of channels auctioned are 10 and 15, respectively. The results show that the privacy leakage follows the same pattern as the case of 5 channels. We can also observe that PASS reduces privacy leakage when the number of channels grows. This is because when the number of channels grows, the number of winning bidders as well as the outcome gets larger, thus preserving privacy better.

These results show that PASS achieves good differential privacy.

VII. Conclusion

In this paper, we have presented a differentially private and strategy-proof spectrum auction mechanism with approximate revenue maximization, which is the first of its kind. Assuming that the seller has prior knowledge of the bidders' valuation distributions, we have first presented a differentially private and strategy-proof auction mechanism which achieves a near optimal expected revenue.To tackle the problem of high computational complexity, we have then presented PASS, an approximate differentially private, strategy-proof, and polynomially tractable auction mechanism. We have theoretically analyzed both mechanisms in terms of strategy-proofness, revenue maximization, and privacy preservation. We have also implemented and evaluated PASS, demonstrating its relatively high revenue generation while preserving bid privacy.

[2]The range of each bidder's valuation/bid, budget, and the value of ε, and δ can be chosen differently from those used here. However, the results of using different setups are similar to those shown in this paper. Therefore, we only show the results for the above setup.

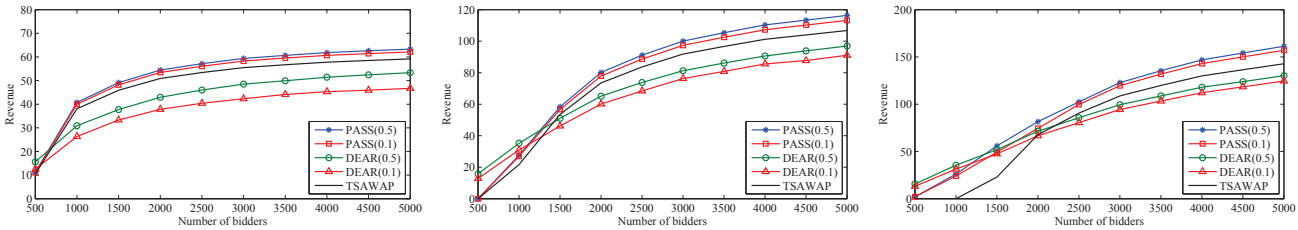

(a) 5 channels (b) 10 channels (c) 15 channels

Fig. 3. Revenue generated by PASS

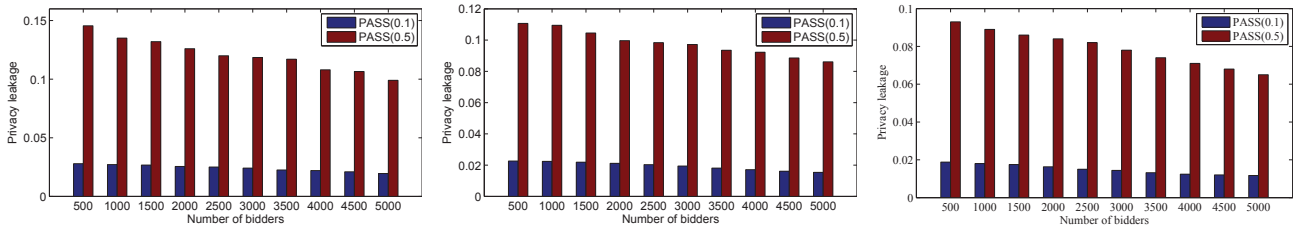

(a) 5 channels (b) 10 channels (c) 15 channels

Fig. 4. Privacy performance

ACKNOWLEDGEMENTS

The work reported in this paper was supported in part by the US Army Research Office (ARO) under Grant W811NF-12-1-0530 and NSF under Grant CNS-1160775.

REFERENCES

[1] M. Al-Ayyoub and H. Gupta, "Truthful spectrum auctions with approximate revenue," in *INFOCOM*, 2011.

[2] A. Archer and E. Tardos, "Truthful mechanisms for one-parameter agents," in *FOCS*, 2001.

[3] Y. Chen, S. Chong, I. A. Kash, T. Moran, and S. Vadhan, "Truthful mechanisms for agents that value privac," in *EC*, 2013.

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

[5] T. Cover and J. Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.

[6] D. C. Cox and D. O. Reudink, "Dynamic channel assignment in high capacity mobile communication system," *Bell System Technical Journal*, vol. 50, no. 6, pp. 1833–1857, 1971.

[7] S. Dobzinski and N. Nisan, "Limitations of vcg-based mechanisms," in *STOC*, 2007.

[8] S. Dobzinski, N. Nisan, and M. Schapira, "Truthful randomized mechanisms for combinatorial auctions," in *STOC*, 2006.

[9] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *TCC*, 2006.

[10] Federal Communications Commission (FCC) Spectrum Auction, http://wireless.fcc.gov/auctions/default.htm?job=auctions_home.

[11] A. V. Goldberg, J. D. Hartline, and A. Wright, "Competitive auctions and digital goods," in *SODA*, 2001.

[12] A. Gupta, K. Ligett, F. McSherry, A. Roth, and K. Talwar, "Differentially private combinatorial optimization," in *SODA*, 2010.

[13] J. Hartline and B. Lucier, "Bayesian algorithmic mechanism design," in *STOC*, 2010.

[14] H. Huang, X.-Y. Li, Y. e Sun, H. Xu, and L. Huang, "Pps: Privacy-preserving strategyproof social-efficient spectrum auction mechanisms," in *http://arxiv.org/pdf/1307.7792v1.pdf*, 2013.

[15] Q. Huang, Y. Tao, and F. Wu, "SPRING: A strategy-proof and privacy preserving spectrum auction mechanism," in *INFOCOM*, 2013.

[16] Z. Huang and S. Kannan, "Exponential mechanism for social welfare: private, truthful, and nearly optimal," in *FOCS*, 2012.

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

[18] G. Kosmopoulou and S. Williams, "The robustness of the independent private value model in bayesian mechanism design," in *Economic theory*, 1998.

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

[20] D. Lehmann, L. I. O'Callaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," in *EC*, 1999.

[21] S. Leung and E. Lui, "Bayesian mechanism design with efficiency, privacy, and approximate truthfulness," in *WINE*, 2012.

[22] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford Press, 1995.

[23] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *FOCS*, 2007.

[24] M. Naor, B. Pinkas, and R. Sumner, "Privacy preserving auctions and mechanism design," in *EC*, 1999.

[25] N. Nisan and A. Ronen, "Algorithmic mechanism design," in *STOC*, 1999.

[26] N. Nisan, "Bidding and allocation in combinatorial auctions," in *EC*, 2000.

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

[28] K. Nissim, C. Orlandi, and R. Smorodinsky, "Privacy-aware mechanism design," in *EC*, 2012.

[29] K. Nissim, R. Smorodinsky, and M. Tennenholtz, "Approximately optimal mechanism design via differential privacy," in *ITCS*, 2012.

[30] A. L. Ny, "Introduction to (generalized) gibbs measures," in *Ensaios Matematicos*, 2008.

[31] M. J. Osborne and A. Rubenstein, *A Course in Game Theory*. MIT Press, 1994.

[32] M. Pan, J. Sun, and Y. Fang, "Purging the back-room dealing: Secure spectrum auction leveraging paillier cryptosystem," in *JSAC*, 2011.

[33] C. H. Papadimitriou, "Algorithms, games, and the internet," in *STOC*, 2001.

[34] D. Wedelin, "An algorithm for large scale $0-1$ integer programming with application to airline crew scheduling," *Annals of Operations Research*, vol. 57, no. 1, pp. 283–301, 1995.

[35] D. Xiao, "Is privacy compatible with truthfulness?" in *ITCS*, 2013.

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

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

[38] R. Zhu, Z. Li, F. Wu, K. G. Shin, and G. Chen, "Differentially private spectrum auction with approximate revenue maximization," in *MobiHoc*, 2014.

[39] R. Zhu and K. G. Shin, "Differentially private and strategy-proof spectrum auction with approximate revenue maximization," in *https://www.dropbox.com/s/z50284bvbgp2vrf/pass.pdf*, 2014.