

Differentially Private Spectrum Auction With Approximate Revenue Maximization

Ruihao Zhu[§] Zhijing Li[‡] Fan Wu[‡] Kang G. Shin[§] Guihai Chen[‡]

[§]Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

[‡]Shanghai Key Laboratory of Scalable Computing and Systems, Shanghai Jiao Tong University, Shanghai, China

Email: {rhzhu,kgshin}@umich.edu; {lizhijing,fwu,gchen}@cs.sjtu.edu.cn

ABSTRACT

Dynamic spectrum redistribution—under which spectrum owners lease out under-utilized spectrum to users for financial gain—is an effective way to improve spectrum utilization. Auction is a natural way to incentivize spectrum owners to share their idle resources. In recent years, a number of strategy-proof auction mechanisms have been proposed to stimulate bidders to truthfully reveal their valuations. However, it has been shown that truthfulness is not a necessary condition for revenue maximization. Furthermore, in most existing spectrum auction mechanisms, bidders may infer the valuations—which are private information—of the other bidders from the auction outcome. In this paper, we propose a **D**ifferentially **p**rivat**E** spectrum auction mechanism with **A**pproximate **R**evue maximization (**DEAR**). We theoretically prove that DEAR achieves approximate truthfulness, privacy preservation, and approximate revenue maximization. Our extensive evaluations show that DEAR achieves good performance in terms of both revenue and privacy preservation.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Wireless Communication

General Terms

Algorithm, Differential Privacy, Economics

Keywords

Spectrum Auction, Mechanism Design, Privacy

1. INTRODUCTION

Radio spectrum is a critical yet scarce resource in this age of fast growing wireless technology. However, the traditional static, expensive, and inefficient spectrum allocation has hampered the growth of the wireless networks and its applications. Under this type of allocation, the utilization of

radio spectrum is low in spatial and temporal dimensions [1]. Thus, a new trend is to make spectrum access dynamic [2]. With dynamic spectrum access (DSA), new wireless applications may take advantage of (instantaneous) spectrum usage opportunities left over by primary spectrum users. Open markets, such as Spectrum Bridge [3], have already emerged to improve spectrum utilization by providing services of selling, buying, and leasing idle spectrum chunks.

Auctions are the best-known market-based allocation mechanisms due to their perceived fairness and efficiency in resource allocation. One important objective of an auction mechanism is to generate a maximum revenue (*i.e.*, total payments collected from the bidders) [8]. Revenue maximization is not only the sellers' goal in the auctions, but also the basic motivation for spectrum owners to sell or lease out under-utilized spectrum.

In recent years, a number of strategy-proof spectrum auction mechanisms (*e.g.*, [10, 11, 12]) with approximate revenue maximization have been proposed, but under a strong assumption that the auctioneer knows the probability distribution of bidders' valuations. However, none of these mechanisms consider privacy preservation. In existing studies, a strategy-proof spectrum auction mechanism motivates each bidder to report her true valuation, which is commonly termed as *type* in the literature. However, once the true valuations of the bidders are reported as bids, one (other than the auctioneer) may infer the private valuation of a bidder based on the outcomes of the auction. In most of the dynamic spectrum auctions, the permissions of using the wireless channels are granted to the bidders for a certain period of time, and the bidders would soon compete again for the usage of the channels. This makes the inference of the bidders' types even easier. For example, n bidders participate in an auction for a set of channels and winners are granted the permissions to use the spectrum within a certain period of time. After the period of time expires, the auctioneer holds another round of the auction for the same set of channels, and the same group of bidders except for one, who leaves the auction, participate in the second round. In both rounds, the participants bid truthfully in order to maximize their utilities, given the strategy-proofness of the auction mechanism. However, the difference between the two auction outcomes may probably expose the type of the bidder who only participates in the first round of the auction to the others. The true types of bidders are critical commercial secrets of the bidders, because they can reflect the potential revenues of the wireless service providers gained through use of the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiHoc'14, August 11–14, 2014, Philadelphia, PA, USA.

Copyright 2014 ACM 978-1-4503-2620-9/14/08 \$15.00.

<http://dx.doi.org/10.1145/2632951.2632974>.

spectrum. Therefore, it is important to achieve bid-privacy preservation in spectrum auctions.

Besides, truthfulness, which means revealing truthful information is the dominant strategy for every bidder, is a too strong assumption, since it prevents the mechanism from having very desirable properties [25, 34]. A number of studies (*e.g.*, [20, 25]) emphasized that proper relaxations on truthfulness can help achieve a good approximation to the revenue and the social welfare in auction mechanisms.

To tackle the problem of privacy preservation, we consider the concept of differential privacy [21], which is a paradigm for private data analysis developed in recent years. Differential privacy aims to reveal information about the population as a whole, while protecting the privacy of each individual. A differentially private mechanism was first investigated by McSherry and Talwar [22]. However, spectrum is different from traditional goods, due to its spatial reusability. Two spectrum users cannot share the same channel simultaneously if their services interfere with each other. Therefore, existing differentially private mechanisms cannot be directly applied to spectrum auctions, in which a slight change of a bidder’s bid may change the outcome dramatically.

In this paper, we aim at designing a differentially private spectrum auction mechanism with approximate revenue without prior information on the distribution of the bidders’ valuations. We present a Differentially privatE spectrum auction mechanism with Approximate Revue (DEAR). In DEAR, the auctioneer will first partition all bidders into groups and subgroups, and then randomly select the payment and allocation which generates approximate revenue with high probability while protecting the privacy.

Our contributions in this paper are summarized as follows.

- Existing differentially private auction mechanisms cannot be directly applied to spectrum allocations, and hence DEAR is the first to bridge differentially private mechanism design and dynamic spectrum redistribution. It protects bidders’ privacy without heavy-weighted cryptographic tools.
- We proposed a polynomial-time spectrum auction mechanism with approximate truthfulness achieving 2ϵ differential privacy and yielding an expected revenue of at least $OPT/7 - 3\ln(e + \epsilon OPT|\mathbb{P}|)/\epsilon$ in a prior-free setting, where OPT is the optimal revenue, \mathbb{P} is the set of possible bids and $\epsilon > 0$ is a small constant.
- DEAR is further extended for the case of multi-channel bids with budget constraints, still guaranteeing 2ϵ differential privacy and providing an expected revenue of at least $OPT/7 - 3c \cdot \ln(e + \epsilon OPT|\mathbb{P}|)/\epsilon$, where c is the number of channels.
- We implement DEAR and extensively evaluated its performance. Our evaluation results show that DEAR generates more revenue than the other mechanisms, while protecting bid-privacy.

The rest of this paper is organized as follows. Section 2 briefly reviews related work, while Section 3 covers technical preliminaries. Section 4 presents the detailed design of DEAR for the single channel request case. Section 5 extends DEAR to support multi-request bidders with budget constraints. Section 6 presents our evaluation results of the proposed mechanism. Finally, we conclude the paper in Section 7.

2. RELATED WORK

In recent years, dynamic spectrum allocation has been suggested as a viable solution to efficiently utilize and share the available spectrum [1, 4, 7]. In this architecture, the spectrum is allocated dynamically in spatial and temporal domains. Periodically, the spectrum owners allocate channels to the bidders under interference constraints with the financial goal.

In traditional economic theory, revenue-maximizing auctions are designed under the assumption that the auctioneer knows the probability distribution of the bidders’ valuations; by applying Vickrey-Clarke-Groves (VCG) mechanism using *virtual valuation* of bidders [13], the resulting auctions can be both strategy-proof and revenue-maximizing (*e.g.*, [10, 11]).

Without prior information of the distribution of bidders’ valuations, Gandhi *et al.* [8] used a linear programming approach to model interference constraints. This work does not consider strategic user behavior, and assumes truthful bids for free. Strategic behavior is considered by Sengupta and Chatterjee, who proposed a knapsack-based auction for dynamic spectrum allocation to optimize the revenue [9]. However, they did not address the problem of interference in spectrum. Considering the spectrum interference constraints, Gopinathan and Li [24] proposed a strategy-proof and revenue maximization spectrum auction mechanism, providing a guarantee on expected revenue. However, none of these mechanisms provides any guarantee on privacy preservation.

The design of privacy-preserving mechanisms has been studied extensively. In [32, 33], the authors presented auction protocols that prune the auctioneers’ ability to falsify the auction outcome and reveal confidential information by introducing a new third party. In [15, 16, 17, 18], the authors employed various cryptography techniques to achieve security in various auction schemes. But these cryptography tools will incur high computation and overheads. Huang *et al.* [14] proposed a strategy-proof and privacy-preserving mechanism called *SPRING*, to achieve k -anonymity in spectrum auctions. Later, Huang *et al.* [36] proposed PPS, which is a privacy-preserving and strategy-proof mechanism for approximate social welfare maximization in spectrum auctions. However, none of these mechanisms considered the problem without prior knowledge of revenue maximization.

Recently, differential privacy was first introduced by Dwork [21]. Then, McSherry and Talwar [22] elegantly integrated differential privacy with mechanism designs, and pointed out that differential privacy implies approximate truthfulness as well as resilience to collusion. In particular, they studied the problem of revenue maximization in digital auctions and attribute auctions. They also suggested use of the exponential mechanism, which can achieve differential privacy, to solve mechanism design problems with different objective functions like revenue. Later studies gave further results on mechanism design via differential privacy, *e.g.*, in [38]. Huang and Kannan examined the properties of the exponential mechanism, which can be thought of as a noisy version of VCG. They showed that, with appropriate payments, this mechanism is truthful, individually rational, approximately efficient, and differentially private.

3. PRELIMINARIES

In this section, we first present the auction model, and then briefly review important solution concepts drawn from differential privacy and mechanism design.

3.1 Auction Model

We consider a collusion-free spectrum auction with an *auctioneer* (seller) and a group of *bidders* (buyers). The auctioneer has a set $\mathbb{C} = \{1, 2, \dots, c\}$ of orthogonal channels. Unlike the allocation of traditional goods, wireless channels can be spatially reused, meaning that multiple well-separated bidders can work on the same channel simultaneously, if they do not interfere with each other. Also, a set $\mathbb{N} = \{1, 2, \dots, n\}$ of bidders compete for the use of channels. Each bidder $i \in \mathbb{N}$ requests a single channel (in Section 4) or multiple channels within its budget (in Section 5). For the case of multi-channel requests, we assume that the bidders do not have preference over different channels.

Each bidder has valuation v_i for a channel. The valuation, which is known as *type* in the literature, is private to the bidder for since can reflect the potential revenue the bidder gains, the valuation of each bidder may need to be kept private. Let $\vec{v} = (v_1, v_2, \dots, v_n)$ denote the profile of the bidders' private valuations. Without loss of generality, we can normalize the bidders' valuations to the interval of $(0, 1]$.

The bidders' bids for the channels are based on their valuations in the auction, and each bidder has a bid b_i per channel. Let $\vec{b} = (b_1, b_2, \dots, b_n)$ denote the profile of the bidders' bids, which are also in the range of $(0, 1]$. The bidders may lie about their valuations, and thus for each bidder i , b_i may not be equal to v_i . The auctioneer also initializes a set of prices $\mathbb{P} = \{\rho_1, \rho_2, \dots, \rho_h\}$, including all the different possible valuation/bid values in $(0, 1]$.

In the auction, the bidders are located in a geographical region. For every bidder, there is a region called *interference area* around her. We say bidders i and j interfere with each other if their interference areas overlap. We assume that a bidder's interference area is a unit-radius disk [11] (*i.e.*, the interference range is 2). The interference among bidders can be model as a *conflict graph* $\mathcal{G} = \{\mathbb{N}, \mathbb{E}\}$, in which the vertices are the bidders, and a pair of bidders is connected if and only if their interference areas overlap.

The auctioneer wishes to allocate channels, such that every pair of the winning bidders do not interfere with each other, and set prices to maximize the total revenue. So, the outcome of an auction includes an allocation profile and a charging profile, which is based on some criteria over the bidding profile \vec{b} . Let $\vec{x} = (x_1, x_2, \dots, x_n)$ denote the allocation profile, where x_i is the number of channels bidder i could get. Let p_i denote the payment for bidder i , and $\vec{p} = (p_1, p_2, \dots, p_n)$ denote the charging profile.

In the auction, each bidder i is considered selfish and rational [29, 30], and always tries to maximize her own utility u_i :

$$u_i = v_i x_i - p_i.$$

The revenue of an auction mechanism is the sum of the payments $\sum_i p_i$ collected from the bidders. Clearly, if the bidders bid untruthfully, the lower bids (than their actual valuations) may indirectly lower the revenue. However, truthfulness is shown to be too strong as a solution, which will impact the allocation scheme and may lower the revenue. So, we want to enforce approximate truthfulness. Specifi-

cally, we aim at maximizing the revenue without any prior information, while enforcing approximate truthfulness and protecting bidders' privacy.

3.2 Related Solution Concept

We now review some of the important and closely related solution concepts from mechanism design and differential privacy.

(1) γ -truthful

We first introduce *Dominant Strategy* [30], a strong solution concept from mechanism design.

DEFINITION 1. (*Dominant Strategy* [30]) *Strategy s_i is a player i 's dominant strategy in a game, if for any strategy $s'_i \neq s_i$ and any other players' strategy profile s_{-i} ,*

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}).$$

The concept of dominant strategy is based on truthfulness. Truthfulness in an auction means that revealing truthful information is a dominant strategy for every bidder. However, exact truthfulness sometimes turns out to be too strict as a solution, so we consider approximate truthfulness, or γ -truthfulness [23].

DEFINITION 2. (γ -truthful [23]) *Let s_i denote the strategy when player i behaves truthfully. A mechanism is said to be γ -truthful if for every player i , for any strategy $s'_i \neq s_i$ and any other players' strategy profile s_{-i} ,*

$$E[u_i(s_i, s_{-i})] \geq E[u_i(s'_i, s_{-i})] - \gamma,$$

where $\gamma > 0$ is a small constant.

(2) Differential Privacy

Differential privacy has been studied extensively in the community of theoretical computer science. It guarantees that the probability distributions of possible outcomes are nearly identical, when the (input) data profiles are nearly identical. Formally,

DEFINITION 3. (*Differential Privacy*[21]) *A mechanism \mathcal{M} gives ϵ differential privacy if for any two data profiles D_1 and D_2 differing in a single element, and all $S \subseteq \text{Range}(\mathcal{M})$,*

$$\Pr[\mathcal{M}(D_1) \in S] \leq \exp(\epsilon) \times \Pr[\mathcal{M}(D_2) \in S],$$

where $\epsilon > 0$ is a small constant.

When integrating differential privacy with an auction mechanism, two neighboring profiles are the bid profiles $\vec{b} = (b_1, b_2, \dots, b_n)$ and \vec{b}' which differs from \vec{b} in one bid (added, removed or changed).

A differentially private mechanism \mathcal{M} can address the concern of personal input leakage. In an auction, any change in a bidder's bid won't bring significant changes to the outcome, and thus, the others cannot infer information of this particular bidder just from the outcomes.

(3) Exponential Mechanism

A powerful tool in the literature of differential privacy is the exponential mechanism proposed by McSherry and Talwar [22].

The exponential mechanism is a general technique for constructing differentially private algorithms over an arbitrary range P of outcomes and any objective function $F(\vec{b}, p)$ that

maps a pair consisting of a data profile \vec{b} and a feasible outcome $p \in P$ to a real-valued score.

In our setting, \vec{b} is the declared valuation (bid) profile, p is every winning bidder's payment per channel. In basic DEAR, the objective function $F(\vec{b}, p)$ is the revenue function; and in extended DEAR, the objective function is the revenue function divided by the number of channels. In a spectrum auction, once the graph of bidders' location is given, the allocation scheme \vec{x} is only depends on parameter \vec{b} and p . Thus, we can denote the revenue function as $Q(\vec{b}, p)$, and $Q(\vec{b}, p) = p \sum_i x_i(\vec{b}, p)$.

Given a range P , a data profile \vec{b} , an objective function F , and a small constant ϵ , the exponential mechanism $Exp(P, \vec{b}, F, \epsilon)$ chooses an outcome p from the range P with probability

$$\Pr[Exp(P, \vec{b}, F, \epsilon) = p] \propto exp\left(\epsilon F(\vec{b}, p)\right),$$

THEOREM 1. *The exponential mechanism gives $2\epsilon\Delta$ differential privacy [22].*

Here, Δ is the Lipschitz constant of the objective function F , i.e., for any two data profiles \vec{b} and \vec{b}' differing in a single element, and for any outcome p , the score $F(\vec{b}, p)$ and $F(\vec{b}', p)$ differs by at most Δ .

Theorem 1 highlights that the exponential mechanism will be the most useful when Δ is small. In DEAR, a small Δ ensures that a single change in the bid of a bidder has a small effect on the outcome.

4. DEAR

We now present DEAR, a differentially private spectrum auction mechanism with approximate revenue maximization.

4.1 Design Rationale

DEAR integrates the exponential mechanism with spectrum auction to achieve both approximate revenue maximization and differential privacy. Basically, DEAR is a single price auction, in which all the winning bidders will be charged with a uniform price. The main idea of DEAR is to randomly select a price from a set of prices with a carefully designed probability distribution. For each price, there are a corresponding allocation and a generated revenue. The probability of a price to be chosen is set to be proportional to its corresponding revenue. Thus, approximate revenue maximization and differential privacy can be achieved. We briefly illustrate the design challenges and our ideas in meeting them.

(1) Payment Selection

In a non-informative-prior setting, the auctioneer has no information on the distribution of bidders' valuations. The obstacle in such a setting is to design a mechanism that can select the optimal/sub-optimal price to charge bidders. Setting p either too high or too low may reduce the revenue. Moreover, we want the payment selection to guarantee approximate truthfulness, which should force bidders in the auction to have limited incentive to lie. To achieve these objectives, we integrate the exponential mechanism with spectrum auction to determine the price. DEAR sets the probability of each price to be chosen to be proportional

to its corresponding generated revenue. Therefore, DEAR could choose the price that generates a high revenue with high probability, while protecting privacy.

(2) Spectrum Allocation

For each price ρ in the set of prices, the auctioneer wants to allocate channels to as many bidders as possible in order to maximize her revenue. Unlike traditional goods, spectrum is reusable among bidders subject to the spatial interference constraints. It has been shown that it's NP-complete [35] to find the optimal spectrum allocation. DEAR partitions the bidders into groups according to their geographic locations. For each price ρ , DEAR allocates channels to groups that can generate the highest revenue.

4.2 Design Details

Following the guidelines in Section 4.1, we now describe the detailed design of DEAR. It performs the auction in three steps. It first partitions the bidders into groups and subgroups. Then, it initializes the set of prices, and calculates the probability distribution over the set. Finally, it randomly selects a price from the set of prices as the payment with the carefully designed probability distribution and allocates the channels to the winners.

Step 1: Grouping of Bidders

Before running the spectrum auction, DEAR first divides the set of bidders \mathbb{N} into groups and subgroups.




Figure 1: Hexagons uniformly-colored using 7 colors. No adjacent hexagons are in the same color.

Since the bidders are located in a geographic region, DEAR basically divides the entire region into small hexagons with unit side-length [11], and then uniformly colors the hexagons with seven colors, as illustrated in Fig. 1.

Let g_i denote the group that contains all the subgroups in color i , and g_i^j denote the j 'th subgroup of g_i which contains the bidders located in the same hexagonal region. Like in Fig. 1, g_1 contains all the bidders located in the yellow region, and g_1^j ($j = 1, \dots, |g_1|$) contains bidders located in each separate yellow hexagon.

Since there are seven colors, we denote the set of groups as:

$$\mathbb{G} = \{g_1, g_2, \dots, g_7\},$$

and denote the set of subgroups in group g_k , $k \in \{1, 2, \dots, 7\}$ as:

$$g_k = \{g_k^1, g_k^2, \dots, g_k^{|g_k|}\}.$$

In such a grouping, we have the following two properties.

Property 1 Any pair of bidders from the same subgroup cannot be allocated on the same channel due to the interference constraint.

Property 2 Any pair of bidders from different subgroups in the same group can share the same channel.

Any pair of bidders from the same subgroup are in the same hexagon with unit side-length (*i.e.*, half of the interference range), so the distance between them is no more than 2. Since every bidder's interference area is a unit-radius disk, then any two bidders' interference areas intersect with each other, so no two bidders from the same subgroup could share a common channel. Thus, Property 1 holds.

Property 2 can be obtained from the observation that since each bidder's interference area is a unit-radius disk, the distance between any pair of bidders from two different subgroups but in the same group is at least $(\frac{\sqrt{3}}{2})^2 + (\frac{5}{2})^2 = \sqrt{7}$. $\sqrt{7} > 2$, which means that the distance between the two bidders is far enough for these bidders to be out of interfere with each other. Thus, they can share the same channel.

These two properties tell us that: (A) For each subgroup g_k^i , the number of winning bidders in it is at most $\min(|g_k^i|, c)$; (B) Winning bidders from different subgroups in the same group can be combined to obtain the whole set of winners.

Step 2: Calculation of Probability Distribution

For each price $\rho_j \in \mathbb{P}$, DEAR removes the bidders with bids less than ρ_j in every subgroup. This changes $g_k^i \in \mathcal{G}$ to $g_k^{i'}$, and changes $g_k \in \mathbb{G}$ to g_k' , and we denote this step as

$$\mathbb{G}' = \text{Remove}(\mathbb{N}, \mathbb{G}, \rho_j),$$

where $\text{Remove}()$ is the removing function that returns the set of groups containing bidders with bids at least ρ_j .

For the remaining bidders in subgroup $g_k^{i'} \in \mathcal{G}'$, DEAR randomly selects $\min(|g_k^{i'}|, c)$ bidders as candidates. Obviously, this method of candidate selection is not related to bidders' bids. We denote this step as

$$W_k^i(\rho_j) = \text{Select}(g_k^{i'}, \min(|g_k^{i'}|, c)),$$

where $\text{Select}()$ is the candidate selection function.

For each group $g_k' \in \mathbb{G}'$, DEAR combines the candidates from all its subgroups to form the set of candidates $W_k(\rho_j)$ in group g_k' at price ρ_j :

$$W_k(\rho_j) = \bigcup_{i=1}^{|g_k'|} W_k^i(\rho_j).$$

DEAR picks up the group with most candidates from the seven groups, let $W(\rho_j)$ be the group with the most candidates when the price is ρ_j :

$$k_0 = \underset{k}{\operatorname{argmax}} |W_k(\rho_j)|, \quad (1)$$

$$W(\rho_j) = W_{k_0}(\rho_j) \quad (2)$$

The tentative price for the candidates in $W(\rho_j)$ is set to ρ_j . Thus, the revenue when setting the price to ρ_j is:

$$Q(\vec{b}, \rho_j) = \rho_j |W(\rho_j)|. \quad (3)$$

DEAR calculates the corresponding revenue when setting the price to all possible values in \mathbb{P} . Then, DEAR sets the probability of price $\rho_j \in \mathbb{P}$ to be chosen proportional to its corresponding revenue, *i.e.*,

$$Pr(\rho_j) = \frac{\exp(\epsilon Q(\vec{b}, \rho_j))}{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i))}.$$

Algorithm 1 Calculation of Probability over Prices

Input: A set of bidders \mathbb{N} , a set of channels \mathbb{C} , a bid profile \vec{b} , a set of prices \mathbb{P} and a set of groups \mathbb{G} .

Output: A probability vector of possible prices \vec{Pr} .

```

1: for all  $\rho_j \in \mathbb{P}$  do
2:    $G' \leftarrow \text{Remove}(\mathbb{N}, \mathbb{G}, \rho_j)$ .
3:   for all  $g_k' \in \mathbb{G}'$  do
4:      $W_k(\rho_j) \leftarrow \emptyset$ .
5:     for all  $g_k^{i'} \in g_k'$  do
6:        $W_k^i(\rho_j) \leftarrow \text{Select}(g_k^{i'}, \min(|g_k^{i'}|, c))$ .
7:        $W_k(\rho_j) \leftarrow W_k(\rho_j) \cup W_k^i(\rho_j)$ .
8:     end for
9:   end for
10:   $k_0 \leftarrow \underset{k}{\operatorname{argmax}} |W_k(\rho_j)|$ .
11:   $W(\rho_j) \leftarrow W_{k_0}(\rho_j)$ 
12:   $Q(\vec{b}, \rho_j) \leftarrow \rho_j |W(\rho_j)|$ .
13: end for
14: for all  $\rho_j \in \mathbb{P}$  do
15:    $Pr(\rho_j) \leftarrow \frac{\exp(\epsilon Q(\vec{b}, \rho_j))}{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i))}$ .
16: end for
Return  $\vec{Pr}$ ;

```

Algorithm 1 shows the pseudo-code of the probability calculation over the prices in set \mathbb{P} .

Step 3: Price Selection and Channel Allocation

After calculating the probabilities of all the prices in set \mathbb{P} to be chosen as payment, we get the probability vector $\vec{Pr} = (Pr(\rho_1), Pr(\rho_2), \dots, Pr(\rho_{|\mathbb{P}|}))$. DEAR randomly selects a price $p \in \mathbb{P}$ as the auction payment per channel according to the probability vector. Then, the corresponding candidates $W(p)$ are the winners and will be allocated channels.

4.3 Analysis

We now analyze the privacy, approximate truthfulness, and revenue of DEAR.

First, we show that the carefully designed grouping can limit the Lipschitz constant Δ of the objective function Q . This way, DEAR can guarantee good differential privacy.

THEOREM 2. *DEAR achieves 2ϵ differential privacy.*

PROOF. We consider two bid profiles \vec{b} and \vec{b}' differing in only one bid. Let M denote the outcome function, which corresponds to the per-channel charge determined by the auctioneer in DEAR. We denote the probability of price $p \in \mathbb{P}$ to be chosen when the bid profile is \vec{b} and \vec{b}' as $Pr(\mathcal{M}(\vec{b}) = \rho)$ and $Pr(\mathcal{M}(\vec{b}') = \rho)$, respectively. We define $\Delta Q \geq 0$, which is the Lipschitz constant of the objective function Q , to be the largest possible difference in Q , when applied to two bid profiles that differ only in one bid, for all p . Then,

$$\begin{aligned}
& \frac{Pr(\mathcal{M}(\vec{b}) = \rho)}{Pr(\mathcal{M}(\vec{b}') = \rho)} \\
&= \frac{\exp(\epsilon Q(\vec{b}, \rho)) / \sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i))}{\exp(\epsilon Q(\vec{b}', \rho)) / \sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}', \rho_i))} \\
&= \frac{\exp(\epsilon Q(\vec{b}, \rho))}{\exp(\epsilon Q(\vec{b}', \rho))} \times \frac{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}', \rho_i))}{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i))}
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{\exp(\epsilon Q(\vec{b}', \rho) + \epsilon \Delta Q)}{\exp(\epsilon Q(\vec{b}', \rho))} \times \frac{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}', \rho_i))}{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i))} &\leq \sum_{(\rho \in \mathbb{P}) \wedge (\rho \leq v_i)} \Pr[M(b_i, b_{-i}) = \rho] \\
&\leq \exp(\epsilon \Delta Q) \times \frac{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i) + \epsilon \Delta Q)}{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i))} &\leq \sum_{(\rho \in \mathbb{P}) \wedge (\rho \leq v_i)} \Pr[M(v_i, b_{-i}) = \rho] \\
&= \exp(\epsilon \Delta Q) \times \exp(\epsilon \Delta Q) &\leq E[u_i(v_i, b_{-i}, v_i, \mathbb{P})] \\
&= \exp(2\epsilon \Delta Q) \tag{4}
\end{aligned}$$

Since in DEAR, for any $\rho \in \mathbb{P}$, any bidder $i \in \mathbb{N}$ can change the revenue function $Q(\vec{b}, \rho)$ by at most

$$\Delta Q = px_i(\vec{b}, \rho) \leq \rho \leq 1. \tag{5}$$

From (1) and (2), we get

$$\Pr(\mathcal{M}(\vec{b}) = p) \leq \exp(2\epsilon) \Pr(\mathcal{M}(\vec{b}') = p). \tag{6}$$

Therefore, DEAR achieves 2ϵ differential privacy. \square

Next, we show the approximate truthfulness of DEAR.

THEOREM 3. *DEAR is 4ϵ -truthful.*

PROOF. Let $E[u_i(b_i, b_{-i})]$ denote the expected utility of bidder i , when she bids b_i . M is the outcome function and ϵ is very small (satisfying $\exp(2\epsilon) < 1 + 4\epsilon$).

Suppose bidder i bids $b_i \neq v_i$. We distinguish the following two cases:

- **Case 1:** We consider the scenario where i bids lower than her true value, $b_i \leq v_i$.

For all $\rho \in \mathbb{P}$, if $b_i < \rho$, the bidder will be removed and then $x_i(b_i, b_{-i}, \rho)v_i - \rho = 0$; if $b_i \geq \rho$, the bidder is still in the subgroup, and $x_i(b_i, b_{-i}, \rho) = x_i(v_i, b_{-i}, \rho)$ for the candidate selection function $Select()$ is not related to the bid. In Theorem 2, we proved that DEAR achieves 2ϵ differential privacy. Obviously, the utility of each bidder with a single demand is less than 1, so we have:

$$\begin{aligned}
&E[u_i(b_i, b_{-i}, v_i, \mathbb{P})] \\
&= \sum_{\rho \in \mathbb{P}} \Pr[M(b_i, b_{-i}) = \rho] \times (x_i(b_i, b_{-i}, \rho)v_i - \rho) \\
&\leq \sum_{\rho \in \mathbb{P}} \exp(2\epsilon) \Pr[M(v_i, b_{-i}) = \rho] \\
&\quad \times (x_i(v_i, b_{-i}, \rho)v_i - \rho) \\
&= \exp(2\epsilon) E[u_i(v_i, b_{-i}, v_i, \mathbb{P})] \\
&\leq (1 + 4\epsilon) E[u_i(v_i, b_{-i}, v_i, \mathbb{P})] \\
&\leq E[u_i(v_i, b_{-i}, v_i, \mathbb{P})] + 4\epsilon \tag{7}
\end{aligned}$$

- **Case 2:** We consider the scenario where $b_i > v_i$.

For all $\rho \in \mathbb{P}$, if $v_i \geq \rho$, then $b_i > \rho$ and bidder i will still be in the subgroup. Since the selection function $Select()$ is not related to the bid, $x_i(b_i, b_{-i}, \rho) = x_i(v_i, b_{-i}, \rho)$. For all $\rho \in \mathbb{P}$ satisfying $v_i < \rho$, if $b_i < \rho$, then bidder i will be removed and the utility is 0; else if $b_i \geq \rho$, bidder i will still be in the subgroup. If she loses, her utility is 0; if she wins, her utility

$$u_i(b_i, b_{-i}, v_i, \rho) = v_i - \rho < 0.$$

Thus, for all $\rho \in \mathbb{P}$ satisfying $v_i < \rho$, misreporting leads the utility to non-positive. Then,

$$\begin{aligned}
&E[u_i(b_i, b_{-i}, v_i, \mathbb{P})] \\
&= \sum_{\rho \in \mathbb{P}} \Pr[M(b_i, b_{-i}) = \rho] \times (x_i(b_i, b_{-i}, \rho)v_i - \rho) \tag{8}
\end{aligned}$$

Thus, the theorem follows. \square

THEOREM 4. *DEAR has an expected revenue of at least $OPT/7 - 3\ln(e + \epsilon OPT|\mathbb{P}|)/\epsilon$, where $|\mathbb{P}|$ is the number of different bids from \mathbb{P} , which is the set of price.*

To prove this theorem, we first give the following lemmas.

Let $OPT^* = \max_{\rho} Q(\vec{b}, \rho) = \max_{\rho} \rho |W_1(\rho)'|$, let OPT denote the optimal single price revenue for the spectrum auction.

LEMMA 1. *OPT^* is within a factor of 7 of OPT , i.e., $OPT/7 \leq OPT^* \leq OPT$.*

PROOF. Suppose the revenue of the spectrum auction reaches OPT when the charged price to each winner is ρ' . Then, for price ρ' , DEAR chooses the group with the most candidates (also with the highest revenue) from the seven groups. Since OPT is upper-bounded by the sum of the optimal revenue of each of the seven groups. Thus, $OPT/7 \leq OPT^* \leq OPT$. \square

LEMMA 2. *Letting $S_t = \{\rho : Q(\vec{b}, \rho) > OPT^* - t\}$. For those t satisfying $t \geq \ln(|\mathbb{P}|OPT^*/(t|S_t|))/\epsilon$, the expected revenue generated by our mechanism $E[REV(\mathcal{M})] \geq OPT^* - 3t$.*

PROOF. Let $\overline{S_{2t}} = \{\rho : Q(\vec{b}, \rho) \leq OPT^* - 2t\}$, we first prove that $\Pr(\rho \in \overline{S_{2t}})$ is at most $\frac{|\mathbb{P}| \exp(-\epsilon t)}{|S_t|}$.

The probability $\Pr(\rho \in \overline{S_{2t}})$ is no more than $\Pr(\rho \in \overline{S_{2t}})/\Pr(\rho \in S_t)$, as the new denominator is at most one. Then, we can write:

$$\begin{aligned}
&\frac{\Pr(\rho \in \overline{S_{2t}})}{\Pr(\rho \in S_t)} \\
&= \frac{\sum_{\rho \in \overline{S_{2t}}} \exp(\epsilon Q(\vec{b}, \rho)) / \sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i))}{\sum_{\rho \in S_t} \exp(\epsilon Q(\vec{b}, \rho)) / \sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i))} \\
&= \frac{\sum_{\rho \in \overline{S_{2t}}} \exp(\epsilon Q(\vec{b}, \rho))}{\sum_{\rho \in S_t} \exp(\epsilon Q(\vec{b}, \rho))} \\
&< \frac{\exp(\epsilon(OPT^* - 2t)) |\overline{S_{2t}}|}{\exp(\epsilon(OPT^* - t)) |S_t|} \\
&= \exp(-\epsilon t) \frac{|\overline{S_{2t}}|}{|S_t|} \\
&\leq \frac{|\mathbb{P}| \exp(-\epsilon t)}{|S_t|}. \tag{9}
\end{aligned}$$

Then, DEAR selects price $\rho \in \mathbb{P}$ that achieves revenue at least $OPT^* - 2t$ with probability of at least $1 - \frac{|\mathbb{P}| \exp(-\epsilon t)}{|S_t|}$. By using the assumption on t , we make this probability at least $1 - \frac{t}{OPT^*}$. Multiplying them, we get:

$$\begin{aligned}
E[REV(\mathcal{M})] &\geq (1 - \frac{t}{OPT^*})(OPT^* - 2t) \\
&> OPT^* - 3t. \tag{10}
\end{aligned}$$

\square

Next, we use the above lemmas to prove Theorem 4.

PROOF. Let $t = \ln(e + \epsilon OPT^* |\mathbb{P}|) / \epsilon$, notice that $t \geq 1/\epsilon$, then we have,

$$\begin{aligned} \ln(OPT^* |\mathbb{P}| / (t |S_t|)) / \epsilon &< \ln(OPT^* |\mathbb{P}| / t) / \epsilon \\ &< \ln(e + OPT^* |\mathbb{P}| / t) / \epsilon \\ &< \ln(e + \epsilon OPT^* |\mathbb{P}|) / \epsilon \end{aligned} \quad (11)$$

We apply Lemma 2 using $t = \ln(e + \epsilon OPT^* |\mathbb{P}|) / \epsilon$ and by Lemma 1, we have:

$$\begin{aligned} E[REV(\mathcal{M})] &\geq OPT^* - 3t \\ &\geq \frac{1}{7} OPT - \frac{3}{\epsilon} \ln(e + \epsilon OPT^* |\mathbb{P}|) \\ &\geq \frac{1}{7} OPT - \frac{3}{\epsilon} \ln(e + \epsilon OPT |\mathbb{P}|). \end{aligned} \quad (12)$$

This completes the proof. \square

5. EXTENSION TO MULTI-DEMAND BIDDERS WITH BUDGET CONSTRAINTS

In the previous section, we proposed a differentially private spectrum auction mechanism with approximate revenue maximization, in which each bidder bids for a single channel. Here we extend it to adapt to a more practical scenario in which a bidder can bid for multiple channels with a budget constraint. This extension also achieves privacy preservation, approximate truthfulness and approximate revenue maximization.

Different from the existing spectrum auction with bidders having multiple requests, we consider the budget of each bidder instead of the demand number. The budget of a bidder is the maximum upper bound on her ability to pay; this is a very common constraint. Also, budget constraints are studied extensively in the theoretical computer science community and are a central feature of many real auctions.

In the extended mechanism, we allow bidders to bid for multiple channels by submitting their budgets. Let $\mathbb{B} = (B_1, B_2, \dots, B_n)$ denote the budget profile of the bidders. We assume that each bidder has an identical valuation on different channels. In the auction, each bidder $i \in \mathbb{N}$ submits not only her per-channel bid $b_i \in (0, 1]$, but also the budget $B_i \in [0, c]$ to the auctioneer. Here, c is the number of channels.

5.1 Virtual Bidder

In the extended DEAR, bidders with adequate budgets may get more than one channels. Here, we introduce the concept of virtual bidder to turn the extended problem into the basic problem.

For price $\rho \in \mathbb{P}$, bidder i with a bid greater than ρ could get at most $\lfloor B_i / \rho \rfloor$ different channels. Thus, DEAR creates $\lfloor B_i / \rho \rfloor$ replicas called *virtual bidders* for bidder i . Each virtual bidder can be allocated at most one channel. DEAR places these virtual bidders in the same subgroup with bidder i . Therefore, any two virtual bidders of a bidder should be allocated different channels. Then, DEAR sets the bids of all the virtual bidders to ρ .

Let's consider an illustrative example of virtual bidder.

Suppose there are 20 channels in the auction, bidder $i \in g_k^j$ is willing to buy spectrum at the price of 0.4 per channel, and she has a budget of 0.9. Assume a price from the price set is $\rho = 0.3$. Given p , bidder i could buy at most $\lfloor 0.9/0.3 \rfloor$




Figure 2: An illustrative example.

$= 3$ different channels and will be charged at 0.3 per channel. So, DEAR creates 3 virtual bidders i^1 , i^2 and i^3 instead of the original bidder i in subgroup g_k^j , and then set the bids of these virtual bidders to 0.3, as shows in Fig. 2.

5.2 Extended DEAR

The procedures of bidder grouping and random selection and allocation are nearly the same as those in DEAR. Due to space limitation, we will focus on the differences in the step of probability calculation.

Step 1: Bidder Grouping

Please refer to subsection 4.2 for details.

Step 2: Calculation of Probability Distribution

For each price $\rho_j \in \mathbb{P}$, DEAR first removes the bidders with bids less than ρ_j in every subgroup. This changes each subgroup $g_k^i \in g_k$ to $g_k^{i'}$, changes each group $g_k \in \mathbb{G}$ to g_k' , and changes \mathbb{G} to \mathbb{G}' . Then, for the remaining bidders in each subgroup $g_k^{i'} \in g_k'$, DEAR creates virtual bidders for each bidder $l \in g_k^{i'}$ according to her budget and the price ρ_j . Let $\hat{g}_k^{i'}$ denote the set of virtual bidders in subgroup $g_k^{i'}$.

DEAR randomly selects $\min(|\hat{g}_k^{i'}|, c)$ virtual bidders as virtual candidates, represented by $W_k^{\hat{g}_k^{i'}}(\rho_j)$. This method of virtual candidate selection is not related to bidders' bids.

For each group $g_k' \in \mathbb{G}'$, DEAR combines the virtual candidates from all its subgroups to form the set of virtual candidates $W_k(\rho_j)$ in group g_k' at price ρ_j :

$$W_k(\rho_j) = \bigcup_{i=1}^{|\hat{g}_k|} W_k^{\hat{g}_k^i}(\rho_j).$$

DEAR picks up the group with most virtual candidates from the seven groups, let $W(\rho_j)$ be the group with the most virtual candidates when the setting price is ρ_j .

The tentative price for the virtual candidates in $W(\rho_j)$ is set to ρ_j . Thus, the revenue when setting the price to ρ_j is:

$$Q(\vec{b}, \rho_j) = \rho_j |W(\rho_j)|.$$

DEAR calculates the corresponding revenue when setting the price to all possible values in \mathbb{P} . Then, DEAR sets the probability of price $\rho_j \in \mathbb{P}$ to be chosen proportional to its corresponding revenue divided by the number of channels, *i.e.*,

$$Pr(\rho_j) = \frac{\exp(\epsilon Q(\vec{b}, \rho_j) / c)}{\sum_{\rho_i \in \mathbb{P}} \exp(\epsilon Q(\vec{b}, \rho_i) / c)}.$$

Step 3: Price Selection and Channel Allocation

DEAR selects the price and virtual winners according to subsection 4.2. Winners will be the bidders of the winning

virtual bidders, and the number of channels each winner gets is equal to the number of her virtual winners.

For the extended DEAR, we get the following theorem.

THEOREM 5. *The extended DEAR is individually rational, and every bidder pays within her budget.*

PROOF. Suppose DEAR selects $p \in \mathbb{P}$ as the payment and allocates channels to the corresponding winning bidders. If the bid of bidder $i \in \mathbb{N}$ is less than p , then she will be charged nothing. Otherwise, bidder $i \in \mathbb{N}$ will get at most $\lfloor B_i/\rho \rfloor$ channels and will pay at most $\rho \lfloor B_i/\rho \rfloor \leq B_i$. Thus, every bidder pays within her budget. The mechanism achieves individual rationality. \square

THEOREM 6. *The extended DEAR gives 2ϵ differential privacy, an expected revenue at least $OPT/7 - 3c \cdot \ln(e + \epsilon OPT/|\mathbb{P}|)/\epsilon$, where $|\mathbb{P}|$ is the number of different bids of \mathbb{P} , and the extended DEAR is $4\epsilon(\max_{i \in \mathbb{N}} E[u_i])$ -truthful.*

PROOF. Here, the objective function is $Q(\vec{b}, \rho)/c$.

As for any $\rho \in \mathbb{P}$, any bidder $i \in \mathbb{N}$ can change the revenue function $Q(\vec{b}, \rho)$ by at most

$$p \lfloor B_i/\rho \rfloor \leq B_i \leq c.$$

Thus, $\Delta(Q/c)$ could be at most 1. Similar to the proof of Theorem 2, the extended mechanism gives 2ϵ differential privacy.

The proof of approximate truthfulness and revenue analysis is similar to the proof of Theorem 3 and Theorem 4 respectively. \square

6. EVALUATION

We have implemented DEAR and extensively evaluated its performance. On one hand, our evaluation results show that DEAR can generate a relatively high revenue despite randomness. On the other hand, the evaluation results show that DEAR achieves good differential privacy.

6.1 Methodology

When evaluating the revenue of DEAR, we compare DEAR with a greedy and truthful spectrum auction mechanism (denoted by ‘‘GREEDY’’) proposed in [10] and another truthful spectrum auction mechanism (denoted by ‘‘TSAWAP’’) in [11]. For the extended DEAR, since there exists no other privacy-preserving spectrum auction mechanism for bidders with budget constraints, we evaluate its performance under various settings without comparison with others.


We vary the number of bidders from 100 to 1500 with a step of 100, and set the number of channels to be 15, 20, and 30. The bidders are randomly deployed in a square area of $5000\text{m} \times 5000\text{m}$. Each bidder has an interference range of 425m [37]. Any pair of bidders who lie within each other’s interference range are in conflict, and thus cannot be allocated on the same channel simultaneously. We assume that each bidder’s bid is uniformly distributed over $(0,1]$ with a precision of 2 decimal places. In the case of extended DEAR, we further assume that each bidder has a budget B_i randomly chosen from $[b_i, c]$, c is the number of channels. We set

the privacy constant ϵ to 0.1 and 0.5.¹ All the results are averaged over 1000 runs.


We use two performance metrics to evaluate DEAR, including revenue and privacy. The revenue is referred to as the sum of charges to the bidders. A mechanism guarantees good privacy if the outcome of probability distribution over prices has a as small as possible change when any bidder unilaterally reports a different bid. We define the notion of *Privacy Leakage* according to the definition of differential privacy to quantitatively measure the privacy guarantee of DEAR.

DEFINITION 4 (PRIVACY LEAKAGE). *Given a mechanism \mathcal{M} , suppose \vec{a} and \vec{a}' are probability distributions over a price set \mathbb{P} for bidding profiles \vec{b} and \vec{b}' , which only differ in a single bid, respectively. The privacy leakage between the two bidding profiles is the maximum of absolute differences between the logarithmic probabilities of the two distributions, i.e.,*


$$\max_{i \in \{1, 2, \dots, |\mathbb{P}|\}} |\ln a_i - \ln a'_i|. \quad (13)$$



(a) 15 Channels



(b) 20 Channels



(c) 30 Channels

Figure 3: Revenue comparison among GREEDY, TSAWAP and DEAR.

¹The range of each bidder’s valuation/bid, budget, and the value of ϵ can be chosen differently from those used here. However, the results of using different setups are similar to the results shown in this paper. Therefore, we only show the results for the above setup.

6.2 Revenue

We first evaluate DEAR and extended DEAR’s performances in terms of revenue.

Fig. 3(a) shows the comparison results of DEAR with GREEDY and TSAWAP, when the number of channels is 15. From these results, we can see that DEAR outperforms GREEDY and TSAWAP in nearly all the cases in terms of revenue. The only exceptions are when the number of bidders is greater than 1200 (for $\epsilon = 0.5$), and when the number of bidders is greater than 1000 (for $\epsilon = 0.1$), GREEDY and TSAWAP can generate slightly more revenues than DEAR. This is because both GREEDY and TSAWAP rely on the existence of critical neighbors to generate revenue, and when the number of bidders grows, GREEDY and TSAWAP can find critical neighbors with higher and higher bids.

Figs. 3(b) and 3(c) show the comparison results of DEAR with GREEDY and TSAWAP, when the number of channels are 20 and 30, respectively. From these results, we can see that DEAR outperforms GREEDY and TSAWAP significantly in all the cases in terms of revenue. This is because DEAR does not rely on the existence of critical neighbor to generate revenue, and when the number of channels grows, some winning bidders may not have any critical neighbor.

These results imply that DEAR is suitable for secondary spectrum markets with relatively sparse bidders.

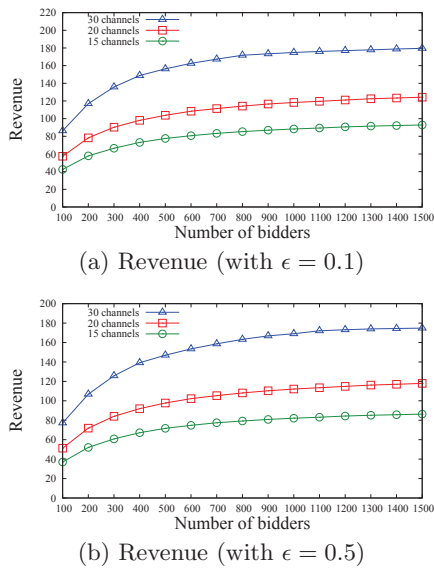


Figure 4: Revenue generated by extended DEAR

Fig. 4 shows the revenue generated by extended DEAR, when bidders bid for multiple channels with budget limits. We observe that, the larger the number of channels is, the higher the revenue is. When the number of channels is fixed, the growth rate of revenue becomes slower as the number of bidders increases. This is because the number of winning bidder in each subgroup is bounded by the number of channels, and thus, the generated revenue converges.

6.3 Privacy

We finally evaluate DEAR and extended DEAR in terms of privacy. We set the number of channels to 20, and vary the number of bidders from 100 to 1500. Fig. 5 shows the privacy leakage of DEAR and extended DEAR.

Fig. 5(a) shows the privacy leakage of DEAR. The results show that when $\epsilon = 0.1$, DEAR’s privacy leakage is always less than 0.018, and when $\epsilon = 0.5$ DEAR’s privacy leakage is always less than 0.085. From the evaluation results, we see that in both cases the privacy performance of DEAR are far better than 0.2–differential privacy and 1–differential privacy, respectively. This implies that it is nearly impossible for any agent to learn the bid information of the others. Thus, DEAR can guarantee good privacy performance.

Fig. 5(b) shows the privacy leakage of extended DEAR. Similar to DEAR, the results show that in both cases of $\epsilon = 0.1$ and 0.5, the privacy leakages of extended DEAR are small, and the privacy performance of DEAR are far better than 0.2–differential privacy and 1–differential privacy, respectively. Besides, we can see that when the number of bidders increases, the privacy leakage of extended DEAR decreases. This is because extended DEAR allows each bidder to bid for multiple channels, and thus the impact of a single bid change is limited.

These results show that (extended) DEAR achieves good differential privacy.

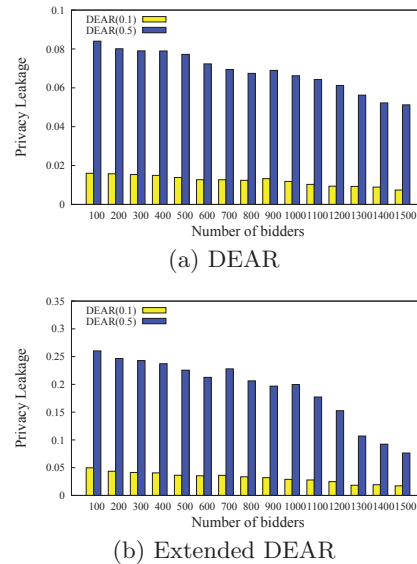


Figure 5: Privacy performance.

7. CONCLUSION

In this paper, we have presented the first differentially private spectrum auction mechanism, called DEAR, with approximate revenue maximization. DEAR performs well with both single- and multi-channel requests when bidders have budget constraints. For both cases, we have theoretically proven the properties in revenue and privacy. We have also implemented DEAR and extensively evaluated its performance. Our mechanisms are shown to be able to generate relatively high and stable revenue, while protecting the bid-privacy.

8. ACKNOWLEDGEMENT

The UM authors were supported in part by the US NSF under Grant No. CNS-1114837. The SJTU authors of this paper were supported in part by the State Key Devel-

opment Program for Basic Research of China (973 project 2014CB340303 and 2012CB316201), China NSF grant 61272443 and 61133006, Shanghai Science and Technology fund 12PJ1404900 and 12ZR1414900, and Program for Changjiang Scholars and Innovative Research Team in University (IRT1158, PC-SIRT) China. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

9. REFERENCES

- [1] I. F. Akyildiz, W. Y. Lee, M. C. Vuran, and S. Mohanty, "NeXt generation/dynamic spectrum access/cognitive radio wireless networks: A survey," *Computer Networks*, vol. 50, no. 13, pp. 2027-2159, 2006.
- [2] Spectrum Policy Task Force, "Spectrum policy task force report," *Federal Communications Commission ET docket*, 02-135, 2002.
- [3] Spectrum Bridge, Inc., <http://www.spectrumbridge.com>
- [4] Z. Ji and K. Liu, "Dynamic spectrum sharing: A game theoretical overview," *IEEE Communications Magazine*, vol. 45, no. 5, pp. 88-94, 2007.
- [5] X. Zhou and H. Zheng, "TRUST: A general framework for truthful double spectrum auctions," in *INFOCOM'09*, 2009.
- [6] F. Wu and N. Vaidya, "SMALL: A strategy-proof mechanism for radio spectrum allocation," in *INFOCOM'11*, 2011.
- [7] J. Peha, "Approaches to spectrum sharing," *IEEE Communications Magazine*, vol. 43, no. 2, pp. 10-12, 2005.
- [8] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri, "A general framework for wireless spectrum auctions," in *DySPAN'07*, 2007.
- [9] S. Sengupta, and M. Chatterjee, "An economic framework for dynamic spectrum access and service pricing," *IEEE/ACM Transactions on Networking*, vol. 17, no. 4, pp. 1200-1213, 2009.
- [10] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *MobiHoc'09*, 2009.
- [11] M. Al-Ayyoub and H. Gupta, "Truthful spectrum auctions with approximate revenue," in *INFOCOM'11*, 2011.
- [12] A. P. Subramanian, M. Al-Ayyoub, H. Gupta, et al. "Near-optimal dynamic spectrum allocation in cellular networks," in *DySPAN'08*, 2008.
- [13] R. Myerson, "Optimal auction design," *Mathematics of operations research*, vol. 6, no. 1, pp. 58-73, 1981.
- [14] Q. Huang, Y. Tao, and F. Wu, "SPRING: A strategy-proof and privacy preserving spectrum auction mechanism," in *INFOCOM'13*, 2013.
- [15] M. Jakobsson and A. Juels, "Mix and match: Secure function evaluation via ciphertexts," in *ASIACRYPT'00*, 2000.
- [16] K. Sako, "An auction protocol which hides bids of losers," in *PKC'00*, 2000.
- [17] K. Peng, C. Boyd, E. Dawson, and K. Viswanathan, "Robust, privacy protecting and publicly verifiable sealed-bid auction," in *ICICS'02*, 2002.
- [18] Y. F. Chung, K. H. Huang, H. H. Lee, F. Lai, and T. S. Chen, "Bidder-anonymous english auction scheme with privacy and public verifiability," *Journal of Systems and Software*, vol. 81, no. 1, pp. 113-119, 2008.
- [19] A. Archer and E. Tardos, "Frugal path mechanisms," in *TALG'07*, 2007
- [20] J. Feigenbaum, and S. Shenker, "Distributed algorithmic mechanism design: Recent results and future directions," September, in *DialM'02*, 2002.
- [21] C. Dwork, "Differential privacy," in *ICALP'06*, 2006
- [22] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *FOCS'07*, 2007.
- [23] A. Gupta, K. Ligett, F. McSherry, A. Roth and K. Talwar. "Differentially private combinatorial optimization," in *SODA'10*, 2010.
- [24] A. Gopinathan and Z. Li, "A prior-free revenue maximizing auction for secondary spectrum access," in *INFOCOM'11*, 2011.
- [25] A. Kothari, DC. Parkes, and S. Suri, "Approximately-strategyproof and tractable multiunit auctions," *Decision Support Systems*, vol. 39, no. 1, pp. 105-121, 2005.
- [26] A. Mu'Alem, and N. Nisan, "Truthful approximation mechanisms for restricted combinatorial auctions," *Games and Economic Behavior*, vol. 64, no. 1, pp. 612-631, 2008.
- [27] K. Nissim, R. Smorodinsky and M. Tennenholtz, "Approximately optimal mechanism design via differential privacy," in *ITCS'12*, 2012.
- [28] C. Dwork, "Differential privacy: A survey of results," in *TAMC'08*, 2008.
- [29] V. Krishna, "Auction theory," Academic Press, 2002.
- [30] N. Nisan, "Algorithmic game theory," Cambridge University Press, 2007.
- [31] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal on Uncertainty, Fuzziness and Knowledge based Systems*, vol. 10, no. 5, pp. 557-570, 2002.
- [32] M. Naor, B. Pinkas, and R. Sumner, "Privacy preserving auctions and mechanism design," in *EC'99*, 1999.
- [33] H. Lipmaa, N. Asokan, V. Niemi, "Secure Vickrey auctions without threshold trust," in *FC'03*, 2003.
- [34] J. Schummer, "Almost-dominant strategy implementation: exchange economies," *Games and Economic Behavior*, vol. 48, no. 1, pp. 154-170, 2004.
- [35] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," in *MobiCom'03*, 2003.
- [36] H. Huang, X. Li, Y. Sun, H. Xu, and L. Huang, "PPS: Privacy-Preserving Strategyproof Social-Efficient Spectrum Auction Mechanisms," in <http://arxiv.org/pdf/1307.7792v1.pdf>, 2013.
- [37] F. Wu and N. Vaidya, "A Strategy-Proof Radio Spectrum Auction Mechanism in Noncooperative Wireless Networks," in *IEEE Transaction on Mobile Computing*, vol. 12, no. 5, pp. 885-894, 2013.
- [38] Z. Huang and S. Kannan, "The exponential mechanism for social welfare: Private, truthful, and nearly optimal," in *FOCS'12*, 2012.