

PADS: Privacy-preserving Auction Design for Allocating Dynamically Priced Cloud Resources

Jinlai Xu*, Balaji Palanisamy*, Yuzhe Tang[†], S.D. Madhu Kumar[‡]

*School of Computing and Information, University of Pittsburgh, Pittsburgh, PA, USA

Email: {jinlai.xu, bpalan}@pitt.edu

[†]Department of EECS, Syracuse University, Syracuse, NY, USA

Email: ytang100@syr.edu

[‡]Department of CSE, National Institute of Technology Calicut, India

Email: madhu@nitc.ac.in

Abstract—With the rapid growth of Cloud Computing technologies, enterprises are increasingly deploying their services in the Cloud. Dynamically priced cloud resources such as the Amazon EC2 Spot Instance provides an efficient mechanism for cloud service providers to trade resources with potential buyers using an auction mechanism. With the dynamically priced cloud resource markets, cloud consumers can buy resources at a significantly lower cost than statically priced cloud resources such as the on-demand instances in Amazon EC2. While dynamically priced cloud resources enable to maximize datacenter resource utilization and minimize cost for the consumers, unfortunately, such auction mechanisms achieve these benefits only at a cost significant of private information leakage. In an auction-based mechanism, the private information includes information on the demands of the consumers that can lead an attacker to understand the current computing requirements of the consumers and perhaps even allow the inference of the workload patterns of the consumers. In this paper, we propose *PADS*, a strategy-proof differentially private auction mechanism that allows cloud providers to privately trade resources with cloud consumers in such a way that individual bidding information of the cloud consumers is not exposed by the auction mechanism. We demonstrate that *PADS* achieves differential privacy and approximate truthfulness guarantees while maintaining good performance in terms of revenue gains and allocation efficiency. We evaluate *PADS* through extensive simulation experiments that demonstrate that in comparison to traditional auction mechanisms, *PADS* achieves relatively high revenues for cloud providers while guaranteeing the privacy of the participating consumers.

I. INTRODUCTION

With the rapid growth of Cloud Computing technologies, enterprises are increasingly deploying their services in the Cloud. The evolution of cloud computing and datacenter-enabled technologies has significantly revolutionized the way in which users and businesses use computing resources today. The cumulative market for cloud computing services is expected to increase to more than 100 billion in 2017 [1]. Dynamically priced cloud resources such as the Amazon EC2 Spot Instance [2] provides an effective mechanism for cloud service providers to trade resources with potential buyers using an auction mechanism. With the dynamically priced cloud resource markets, cloud consumers can buy resources at a cost much lower than statically priced cloud resources such as the on-demand instances in Amazon EC2 [3]. Such spot instances can reduce the cost of the computing resources by up to 50%

to 90% if the applications running on the spot instances can deal with temporary interruptions during job execution [4]. Thus, Spot Instances are highly recommended for applications such as data mining and batch processing that do not have a real-time processing requirement [4].

While dynamically priced cloud resources enable to maximize datacenter resource utilization and minimize cost for the consumers, unfortunately, such auction mechanisms achieve these benefits only at a significant cost of private information leakage. In an auction-based mechanism, private information includes information on when and who has higher demands on which types of Virtual Machine (VM). Such information can lead an attacker to understand the current computing requirements of the consumers and perhaps even allow the inference of the workload patterns of the consumers. For instance, if a consumer makes a higher bid for the spot instance, an adversary may be able to infer that the requested resources for the computing task are more important than the other resources requested by the user. Such adversaries may also infer other business secrets by combining the bidding information with other background knowledge and break the normal order through false-name bids [5] in the auction market to disrupt the normal fair operation.

Protecting consumer privacy in an auction-based resource allocation market is an important task. Earlier works have addressed how to protect privacy in auctions [6] [7] [8] [9] such that the auction achieves the desired outcomes without revealing the private information of the bidders. While there has been work on privacy-aware auctions in stock and spectrum distribution [10] [11] [12] [13] [14] [15], privacy-preserving auction design for dynamically priced cloud resource allocation has not yet received attention from the research community. In this paper, we propose a privacy-preserving auction design mechanism called *PADS* (privacy-preserving auction design for spot and dynamically priced cloud resources) that protects the private information in the bids in the auction process through differential privacy [16] guarantees.

Concretely, this paper makes the following contributions.

- 1) To the best of our knowledge, the work presented in this paper is the first to design a differentially private and

strategy-proof solution for allocating dynamically priced cloud resources through an auction mechanism.

- 2) We formally model and analyze the problem of dynamically priced cloud resource allocation as a sealed-bid auction problem and design two near optimal privacy-preserving mechanisms. We demonstrate that both the mechanisms achieve differential privacy guarantees and hold the strategy-proof property.
- 3) We propose *PADS-ADP*, an (ϵ, δ) -differentially private and truthful auction mechanism. Unlike existing solutions, *PADS-ADP* has the ability to simultaneously guarantee differential privacy and yet provide the desired features of an auction mechanism. We improve the performance of *PADS-ADP* by developing an enhanced mechanism called *PADS-DP* which is an ϵ -differentially private and an approximate truthful mechanism. The low computational complexity of both *PADS-ADP* and *PADS-DP* make it possible to calculate the auction outcome for low latency real-time scheduling requests.
- 4) We experimentally evaluate *PADS* (both *PADS-ADP* and *PADS-DP*) through an extensive simulation study. Our evaluation results show that *PADS* achieves a closely similar performance compared to traditional auction mechanisms such as VCG [17]–[19] while providing the desired differential privacy guarantees in the auction process.

The remainder of this paper is organized as follows. In Section II, we briefly review the related work in the areas of auction mechanism design and differential privacy. Section III presents the problem model and reviews a few key concepts related to auction design and differential privacy. Section IV presents the design of a near optimal privacy-preserving mechanism and its properties. In Section V and VI, we introduce the design of *PADS-ADP* and *PADS-DP* and analyze their properties. Section VII presents our evaluation results. Finally, we conclude the paper in Section VIII.

II. RELATED WORK

In recent years, privacy becomes a significant concern to users as innovations in technology often require private information of users for processing. There are two kinds of privacy-preserving techniques studied extensively in the literature: k-anonymity [20] and differential privacy [16]. Many privacy-preserving solutions are proposed based on the differential privacy concept, which include two major techniques: the first set of techniques is represented by perturbation algorithms [21], [22] which adds noise to protect the privacy for individual’s contribution to the statistical output; the second set of techniques include exponential mechanisms [9], [23] that is used when the noise addition is not a reasonable approach for guaranteeing privacy. Privacy-preserving auctions also have been studied extensively in the recent years, and some of the work has studied the problem in the context of using differential privacy. For the theoretical aspect, there are many research efforts trying to add new properties to the previous basic auction mechanisms to make them more efficient and effective, for example, McSherry and Talwar




Figure 1. Spot Instance Illustration (Spot Instance Type: g2.8xlarge, Bid Price: \$1.95, Date: June 30, 2017)

[9] presents the basic idea of protecting the privacy of the bids using an exponential mechanism and demonstrate several critical properties including the approximate truthfulness and differential privacy guarantees of the mechanism. Huang and Kannan [23] proposed a nearly optimal differentially private auction mechanism using the Gibbs Measure which is also known as Boltzmann distribution in chemistry and physics [24] to achieve a more effective optimization for the revenue and social welfare compared to the techniques presented in [9]. For the application perspective, spectrum allocation auctions are the most studied ones in the area of privacy-preserving auctions. There are several research efforts implementing privacy preserving mechanisms in spectrum auctions such as [25], [26]. These techniques primarily address the spectrum resource allocation problem in a privacy-preserving manner. To the best of our knowledge, there is no prior work addressing the privacy leakage problem in dynamically priced resource allocation in clouds. The *PADS* privacy-preserving auction mechanism proposed in this work is the first significant effort towards addressing the privacy-preserving cloud resource allocation problem using privacy-aware auctions that provide both differential privacy guarantees and efficiency in terms of resource allocation.

III. CONCEPTS AND MODEL

In this section, we present the problem description for privacy-preserving auction design of the dynamically priced resource allocation and introduce the basic concepts of mechanism design and differential privacy.

A. Problem Model

We model the dynamically priced cloud resource allocation problem as a sequence of auctions using discrete time slots. In each time slot, we assume an auction mechanism that determines the winning bids and allocates the resources to the winners. As shown in Figure 1 which illustrates an example of Spot Instance, when the bid of a user is higher than the spot price, the user can run his/her job using a set of virtual machines (VMs) during the time slot. In the cloud resource allocation auction, we refer to the users as *bidders* or *buyers*

and the Cloud Service Providers (CSPs) as the *sellers*. The entity performing the auction mechanism is referred to as the *auctioneer*. There are K types of VMs used as the goods in the auctions. In every time slot t , for each type- k VM, the sealed-bid auction mechanism decides the users who can run their jobs during the time slot t . For simplicity, we model each round of auction in a time slot assuming only one type of VMs used as goods in the auction. The objective of the seller (CSP) is to allocate the VMs to the users such that it maximizes its revenue. We assume that the CSP (seller) has m type- k VMs. There are n users that want to use the type- k VM and each user $i \in N$ bids for the VMs with their bid value, b_i . Here N denotes the set of bidders $N = \{1, 2, \dots, n\}$. The bids are represented by a vector $\vec{b} = \{b_1, b_2, \dots, b_n\}$. Each user has a per-VM valuation, which is private to the user, represented by $\vec{v} = \{v_1, v_2, \dots, v_n\}$. Depending on the bidding strategy, the bid may be equal or not equal to the real valuation of the good for the user. The outcome of the auction is determined by the auction mechanism which can be represented by a vector $\vec{x} = \{x_1, x_2, \dots, x_n\}$ where x_i is a binary indicator that indicates whether the bid b_i wins or not. The payments are represented by $\vec{p} = \{p_1, p_2, \dots, p_n\}$ where p_i is the payment of user i to rent the type- k VM in the current time slot. The objective of each user is to maximize the per-user utility which can be represented by using the following utility function:

$$u_i = (v_i - p_i)x_i \quad (1)$$

where u_i is the utility of user i . The seller (CSP) also wants to maximize its revenue in the auction mechanism. The revenue of the seller (CSP) can be represented by the sum of the payments:

$$REV = \sum_i^n p_i x_i \quad (2)$$

In a privacy-preserving auction mechanism, one of the objectives is to protect the inference of the participation of a bidder from the outcome of the auction. In addition, the inference of private information such as the bid value, b_i and the user's true valuations of the goods, v_i need to be protected from the outcome of the auction as well.

B. Auction Design Concepts

Before introducing the proposed *PADS* auction mechanism, we review some important concepts related to auction mechanism designs and privacy-preserving mechanisms.

Mechanism Design

We first introduce Dominant Strategy [27] from game theory that forms a fundamental solution concept for auction mechanism designs.

Definition 1. (Dominant Strategy [28]) 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} ,

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}). \quad (3)$$

The concept of dominant strategy is related to truthfulness. In an auction, truthfulness means that revealing truthful information is the dominant strategy for every bidder.

Definition 2. (Truthfulness) "Truthfulness" is also called as strategy-proof or incentive compatibility in auction literature. In game theory, an asymmetric game where players have private information is said to be strategy-proof (SP) if it is a weakly-dominant strategy for every player to reveal his/her private information.

If an auction mechanism is truthful, then the bidders will tend to bid with their true valuation of the products. This is a powerful feature for auction mechanism design as it ensures that both the buyers and sellers can get maximum utility from the auction without cheating.

Formally, we can define the truthfulness property as

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

where the u_i is the utility of bidder i , s_i is the strategy that bidder i bids with the true value of the product. Here s_{-i} represents the strategies for the bidders other than bidder i and s'_i represents a strategy other than s_i . The function illustrates that the strategy that bids with the true value will give the bidder the highest utility compared to any other strategies. If this function is true for all the bidders, it ensures that the auction mechanism is truthful.

However, exact truthfulness sometimes turns out to be too strict as a solution, and as an alternative, approximate truthfulness, or γ -truthfulness [9], [29] has been proposed in the literature.

Definition 3. (γ -truthfulness) An auction is γ -truthful in expectation, or γ -truthful for short, if and only if for any bidding strategy $s'_i \neq s_i$ and for any bid strategies of other bidders s_{-i} , there is:

$$E[u_i(s_i, s_{-i})] \geq E[u_i(s'_i, s_{-i})] - \gamma \quad (4)$$

where γ is a small positive constant.

In auction mechanism design, there is another property called Individual Rationality guaranteeing that every bidder will not lose utility in the auction. It is defined as below:

Definition 4. (Individual Rationality) An auction is individual rational if and only if $u_i \geq 0$ holds for every bidder $i \in N$.

Privacy Concepts and Definitions

We next introduce the concepts and definitions related to privacy-preserving mechanism designs.

Definition 5. (Differential Privacy) Differential privacy is a privacy-preserving mechanism that protects an individual user's contribution in a dataset. In a differentially private auction mechanism, the actions of a trusted auctioneer can be modeled as a randomized algorithm \mathcal{A} . A randomized algorithm \mathcal{A} is ϵ -differentially private if for all datasets D_1

and D_2 that differ on a single element (i.e., a bid of one person), and all $S \subseteq \text{Range}(\mathcal{A})$:

$$\Pr[\mathcal{A}(D_1) \in S] \leq e^\epsilon \times \Pr[\mathcal{A}(D_2) \in S]$$

In the literature, a relaxed definition of differential privacy has also been introduced.

Definition 6. (Approximate Differential Privacy [21]) A randomized algorithm \mathcal{A} is (ϵ, δ) -differentially private if for all datasets D_1 and D_2 that differ on a single element (i.e., a bid of one person), and all $S \subseteq \text{Range}(\mathcal{A})$:

$$\Pr[\mathcal{A}(D_1) \in S] \leq e^\epsilon \times \Pr[\mathcal{A}(D_2) \in S] + \delta$$

Exponential Mechanism is a key mechanism to achieve differential privacy in a privacy-preserving auction design. An exponential mechanism builds a distribution of probabilities to choose the output based on an exponential function to guarantee ϵ -differential privacy [9].

Exponential Mechanism

The exponential mechanism [9] is a general technique for constructing differentially private algorithms over an arbitrary range \mathcal{R} of outcomes and for any objective function $F(b, r)$. The goal of the exponential mechanism is to map, randomly, a set of n inputs each from a domain \mathcal{D} to some output in a range \mathcal{R} and protect the individual privacy. It is defined as follows:

Definition 7. (Exponential Mechanism [9]) For any function $F : (\mathcal{D}^n \times \mathcal{R}) \rightarrow \mathbb{R}$, and a base measure μ over \mathcal{R} , define:

$$\mathcal{E}_F^\epsilon(b) := \text{Choose } r \text{ with probability } \propto e^{\epsilon F(b,r)} \times \mu(r) \quad (5)$$

$, b \in \mathcal{D}^n, r \in \mathcal{R}$

The exponential mechanism guarantees $2\epsilon\Delta$ -differential privacy, where $\Delta = \bar{b} - \underline{b}$ is an upper-bound of the difference of the feasible outcomes of two data sets which only differ in a single data item. An immediate theorem can also be derived as [29]:

Theorem 1. When used to select an output $r \in \mathcal{R}$, the exponential mechanism $\mathcal{E}_F^\epsilon(b)$ yields $2\epsilon\Delta$ differential privacy. Let R_{OPT} denote the subset of \mathcal{R} achieving $F(b, r) = \max_r F(b, r)$, then the exponential mechanism ensures that:

$$\Pr[F(b, \mathcal{E}_F^\epsilon(b)) < \max_r F(b, r) - \ln\left(\frac{|R|}{|R_{OPT}|}\right)/\epsilon - \frac{t}{\epsilon}] \leq e^{-t} \quad (6)$$

In our problem, we want to design an auction mechanism that can allocate the VMs to the users based on the bids submitted by the users and achieve the strategy-proof property while preserving privacy and maximizing the CSP revenue. In the next section, we propose a near optimum mechanism for solving the privacy-preserving auction design problem for cloud resource allocation.

IV. STRAIGHT-FORWARD EXPONENTIAL MECHANISM

We first propose a near optimum privacy preserving mechanism which is straight-forward based on the exponential mechanism proposed in [9] and [23]. The straight-forward mechanism solves the privacy-preserving problem for the dynamically priced resource allocation problem using the exponential mechanism proposed in [9] and [23]. In an auction mechanism for allocating the VMs, \bar{b} represents the bid profile, and p_i represents the bidder i 's payment. The objective for the auction is to maximize the revenue of the CSP which can be calculated as:

Objective:

$$\max REV = \sum_i^n p_i x_i \quad (7)$$

Subject to:

$$\sum_i^n x_i \leq m \quad (8)$$

As the exponential mechanism proposed in [23] achieves approximate truthfulness, the expected revenue is equal to its expected surplus:

$$\mathbf{E}[REV] = \mathbf{E}\left[\sum_i^n v_i x_i\right] \quad (9)$$

The logic behind it is intuitive. In a truthful auction that the bidders tend to bid with the true valuations, the expected revenue of the auction is equal to the expected surplus.

With the above objective function and the expected surplus analysis, the problem of designing the privacy-preserving and near optimum resource allocation auction can be solved through an exponential mechanism. Based on the principles outlined in [23], the privacy-preserving auction mechanism may assign each possible outcome a probability which is proportional to the objective function and can be represented by the revenue $F(\bar{b}, \vec{x}) = \sum_i^n b_i x_i$. Then, based on the probabilities, the mechanism can choose the outcome. The payment for each winner is assigned using a VCG-like mechanism [23].

The detailed auction works as below:

- 1) Each bidder i submits its bid: b_i ;
- 2) The mechanism chooses the outcome with probability proportional to e raised to the power of the objective function and satisfies the Eq. (8):

$$\Pr[\vec{x}] \propto \exp\left(\frac{\epsilon}{\Delta} \sum_i^n b_i x_i\right) \quad (10)$$

- 3) The payment for winner bid b_i is assigned by the mechanism proposed in [23].

A. Analysis

Next, we analyze the features of the near optimum privacy-preserving resource allocation mechanism. In particular, we analyze the revenue guarantee of the mechanism, the truthfulness property and the tractability of the mechanism. As we know that the expected revenue is equivalent to the expected surplus, here we just analyze the expected revenue.

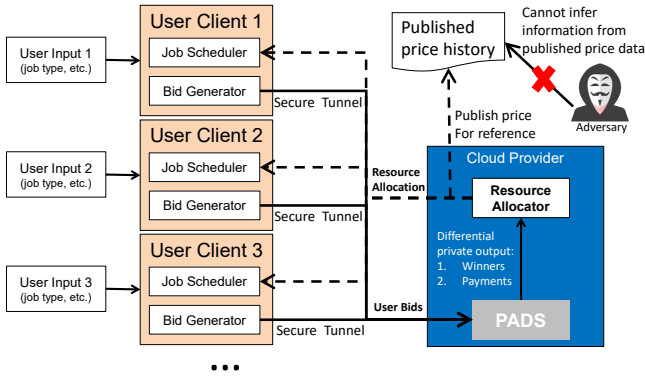


Figure 2. PADS Architecture

Theorem 2. The expected surplus:

$$\mathbf{E}\left[\sum_i^n v_i x_i\right] \quad (11)$$

is maximized when the winning bids are chosen from the near optimum privacy-preserving mechanism, $\mathcal{E}_F^{\epsilon}(\vec{b})$.

Theorem 3. The near optimum privacy-preserving mechanism is truthful, individually rational and ϵ -differentially private.

The proofs of the above-mentioned properties are discussed in [23].

Theorem 4. The exponential mechanism which chooses the outcome according to the objective function to decide the winners for allocating the VMs is intractable.

Proof. The choosing function needs to calculate all the possible outcomes with the constraint defined in the Eq. (8). The computation complexity can be represented by a combination which chooses m from n , $O(\binom{n}{m})$. As the number of VMs (which is represented by m) is usually a large number and the number of users (which is represented by n) is also large, it makes the straight-forward near optimum privacy-preserving mechanism intractable in the dynamically priced resource allocation problem. \square

V. PADS-ADP: PRIVACY-PRESERVING AUCTION DESIGN WITH APPROXIMATE DIFFERENTIAL PRIVACY GUARANTEES

The straight-forward near optimum privacy-preserving mechanism introduced in the previous section can provide near optimum revenue and ϵ -differential privacy. However, the computation cost is prohibitively expensive to be used in practice. In this section, we propose *PADS-ADP*, an alternate privacy-preserving auction design that uses an iterative winner decision algorithm and a payment scheme that forces the bidders to bid with true valuations. We prove that *PADS-ADP* can provide (ϵ, δ) -differential privacy while achieving the truthfulness property in the auction.

In *PADS*, the auctions are conducted in discrete time slots as described in Section III. As shown in Figure 2, similar to

the Spot Instances provided by Amazon EC2, *PADS* assumes that each user has a client which takes care of the bidding and job scheduling. The user submits the maximum price he/she wants to pay for the VMs. The user client bids for the required VMs in every time slot with the maximum price set by the user. When the bid wins in a time slot, the user client schedules the jobs to the VMs allocated to it. As shown in Figure 2 *PADS* protects the private information of the users in the auction including who are the winners and how much they bid and pay for the resources from the adversaries. In an auction performed using *PADS* to determine the winners and payments, the adversaries cannot infer users' information from the published price data (provided by the CSP for the potential customers to refer).

A. Design Details

We now describe the detailed design of *PADS-ADP*. The iterative auction mechanism works in a sequence of four steps: (i) it first calculates the probability distribution over the set of current bids, R , (ii) then it randomly selects a bid from the set as the winner in the current round based on the probabilities calculated from the first step, (iii) next, it calculates the payment scheme for the winner and (iv) finally, it removes the winner from the set, R , in the current round and checks the end condition. The above four steps repeat until there are no bids in the set R or the VMs in the resource pool has been exhausted.

The winners of the auction are determined as follows.

(i) Calculation of Probability Distribution: First, we need to calculate the probability distribution of the bids which needs to be used in the exponential mechanism. The difference between the near optimum solution and *PADS-ADP* is that instead of choosing the results from the all possible outcomes, *PADS-ADP* chooses one winner for each iteration proportional to e raised to the power of the bid value:

$$Pr[W \leftarrow W \cup \{i\}]_i = \frac{\exp(\epsilon' b_i)}{\sum_{i \in R} \exp(\epsilon' b_i)} \quad (12)$$

where W is the set of winners' bids such as $W = \{w_1, w_2, \dots\}$, R is the current set of the bid, and:

$$\epsilon' = \frac{\epsilon}{(e-1)\Delta \ln(e/\delta)} \quad (13)$$

(ii) Winner Selection: After calculating the probabilities of all the bids in set R to be chosen as winners, we get the probability vector $\vec{Pr} = \{Pr_1, Pr_2, \dots\}$. *PADS-ADP* randomly selects a bid $b_i \in R$ as the winner in the current round according to the probabilities for each bid, $Pr_i, \forall i \in R$.

(iii) Payment Scheme: After selecting the winner for the current iteration, *PADS-ADP* calculates the payment for it. As we desire the *truthfulness* property from the auction mechanism, the payment scheme is quite important to make the auction *truthful*. Here, we use the results developed in [30] to design the payment scheme. The immediate theorem can be described as below:

Theorem 5. A mechanism is truthful in expectation if and only if, for any bidder i and any fixed choice of bids by the other bidders, \vec{b}_{-i} ,

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

From the Theorem 5, for the first condition, as we already conduct the exponential mechanism in the winner selection step, the probability of choosing the bidder i is proportional to $\exp(\epsilon' b_i)$. Here, ϵ' is a constant which is a positive number, and the exponential function $\exp()$ is monotonically increasing. Hence the first condition is satisfied. Next, we set the payment scheme as follows:

$$p_i = b_i y_i(\vec{b}) - \int_0^{b_i} y_i(z) dz \quad (14)$$

to satisfy the second condition. Thus, the payment scheme and the above winner selection algorithm together provide *truthfulness* for the auction.

(iv) Post Processing: After calculating the payments, the indicator of the winning bid b_i is set to 1, $x_i = 1$. And the bid, b_i , is removed from the set of current bids, R . Then, *PADS-ADP* checks whether all the VMs are allocated, or all the bids are removed from R . If any of the above two conditions are met, the auction is ended.

Algorithm 1: PADS-ADP Mechanism

Input : Type of the VM : k ;
of VMs: n ;
Buy bids: $\vec{b} = \{b_1, b_2, \dots\}$;
Output: Auction decision: $\vec{x} = \{x_1, x_2, \dots\}$;
Payment scheme: $\vec{p} = \{p_1, p_2, \dots\}$;

- 1 Initially, the possibility vector as $\vec{Pr} = \{Pr_1, Pr_2, \dots\}$ where Pr_i is set by Eq. (12), $R = \vec{b}$ and the number of winners $m = 0$;
- 2 **while** $R \neq \emptyset$ or $m < n$ **do**
- 3 **for** all $b_i \in R$ **do**
- 4 $Pr_i = \frac{\exp(\epsilon' b_i)}{\sum_{i \in R} \exp(\epsilon' b_i)}$;
- 5 **end**
- 6 Randomly select i according to the the probability vector \vec{Pr} ;
- 7 Set $x_i = 1$;
- 8 $p_i = b_i y_i(\vec{b}) - \int_0^{b_i} y_i(z) dz$ Remove b_i from R ;
- 9 $m = m + 1$;
- 10 **end**

The overall algorithm is shown in Algorithm 1. The time complexity of Algorithm 1 is $O(n * m)$ where n is the number of users that bid for the VMs and m is the number of VMs that can be allocated. The worst case time complexity can be estimated by the accumulation of the computation times of the probabilities, \vec{Pr} . It can be calculated as: $n + (n - 1) + \dots + (n - m + 1) = \frac{(n+n-m+1)*m}{2} = n * m + \frac{m^2 + m}{2}$. Generally, the number of users is larger than the number of VMs: $n > m$. Therefore the worst case time complexity of the algorithm is $O(n * m)$.

After proposing the mechanism, we analyze and prove the differential privacy guarantee for *PADS-ADP*.

Theorem 6. For any $\delta \leq 1/2$, *PADS-ADP* provides (ϵ, δ) -differential private.

Proof. Let \vec{b} and \vec{b}' be two input bid vectors that differ in a single bidder s 's bid. We show that *PADS-ADP* achieves bid privacy preservation including revealing the order in which the bidders are chosen for an arbitrary sequence of winners selection $W = W' = \{w_1, w_2, \dots, w_l\}$ of arbitrary length l for \vec{b} and \vec{b}' respectively. The steps of the proof are inspired by the results presented in [25]. We consider the relative probability of *PADS* results for given bids inputs \vec{b} and \vec{b}' :

$$\begin{aligned} & \frac{Pr[W = \{w_1, w_2, \dots, w_l\}]}{Pr[W' = \{w_1, w_2, \dots, w_l\}]} \\ &= \prod_{i=1}^l \frac{\exp(\epsilon' b_{w_i}) / \sum_{j \in (N - \{\pi_i\})} \exp(\epsilon' b_j)}{\exp(\epsilon' b'_{w_i}) / \sum_{j \in (N - \{\pi_i\})} \exp(\epsilon' b'_j)} \quad (15) \\ &= \prod_{i=1}^l \frac{\exp(\epsilon' b_{w_i})}{\exp(\epsilon' b'_{w_i})} \prod_{i=1}^l \frac{\sum_{j \in (N - \{\pi_i\})} \exp(\epsilon' b'_j)}{\sum_{j \in (N - \{\pi_i\})} \exp(\epsilon' b_j)} \end{aligned}$$

where $\pi_i = \{w_1, w_2, \dots, w_i\}$. If $b_s > b'_s$, the first product is less than $\exp(\epsilon' \Delta)$:

$$\exp(\epsilon'(b_s - b'_s)) \leq \exp(\epsilon' \Delta) \quad (16)$$

and the second product is less than 1. If $b_s < b'_s$, the first product is less than 1. Then, we have

$$\begin{aligned} & \frac{Pr[W = \{w_1, w_2, \dots, w_l\}]}{Pr[W' = \{w_1, w_2, \dots, w_l\}]} \\ & \leq \prod_{i=1}^l \frac{\sum_{j \in (N - \{\pi_i\})} \exp(\epsilon' b'_j)}{\sum_{j \in (N - \{\pi_i\})} \exp(\epsilon' b_j)} \quad (17) \\ & = \prod_{i=1}^l \frac{\sum_{j \in (N - \{\pi_i\})} \exp(\epsilon'(b'_j - b_j)) \exp(\epsilon' b_j)}{\sum_{j \in (N - \{\pi_i\})} \exp(\epsilon' b_j)} \\ & = \prod_{i=1}^l \mathbf{E}_{j \in (N - \{\pi_i\})} [\exp(\epsilon'(b'_j - b_j))] \end{aligned}$$

Note that for all $\beta \leq 1$, we have $\exp(\beta) \leq 1 + (e - 1)\beta$. Therefore, for all $\epsilon' \leq 1$, we have

$$\begin{aligned} & \prod_{i=1}^l \mathbf{E}_{j \in (N - \{\pi_i\})} [\exp(\epsilon'(b'_j - b_j))] \\ & \leq \prod_{i=1}^l \mathbf{E}_{j \in (N - \{\pi_i\})} [1 + (e - 1)\epsilon'(b'_j - b_j)] \quad (18) \\ & \leq \exp((e - 1)\epsilon' \sum_{i=1}^l \mathbf{E}_{j \in (N - \{\pi_i\})} [b'_j - b_j]) \end{aligned}$$

If $\sum_{i=1}^l \mathbf{E}_{j \in (N - \{\pi_i\})} [b'_j - b_j] \leq \Delta \ln(e/\delta)$, we have

$$\frac{Pr[W = \{w_1, w_2, \dots, w_l\}]}{Pr[W' = \{w_1, w_2, \dots, w_l\}]} \leq \exp((e - 1)\epsilon' \Delta \ln(e/\delta)) = \exp(\epsilon) \quad (19)$$

By Lemma A.1 and A.2 in [29], we have $Pr[\sum_{i=1}^l \mathbf{E}_{j \in (N - \{\pi_i\})} [b'_j - b_j] > \Delta \ln(e/\delta)] \leq \delta$.

Thus, the theorem follows. \square

VI. PADS-DP: PRIVACY-PRESERVING AUCTION DESIGN WITH DIFFERENTIAL PRIVACY

In the previous section, we presented *PADS-ADP* that can provide (ϵ, δ) -differential privacy and truthfulness. A limiting constraint of *PADS-ADP* is that it can only provide approximate (ϵ, δ) -differential privacy which is relatively weaker than the rigorous ϵ -differential privacy. In addition, as the iterative winner selection makes the mechanism more random with a smaller ϵ' in each iteration, the mechanism is too random to be efficient as demonstrated by our results in the evaluation Section VII. Another weakness of *PADS-ADP* is that the payment scheme is heterogeneous for the winners as it uses a payment scheme which calculates the payments by the probability distributions for each bidder to force the mechanism to be truthful.

In this section, we propose *PADS-DP* which makes a trade-off between differential privacy, the truthfulness property and the revenue earned by the CSP. *PADS-DP* can provide ϵ -differential privacy with little loss of truthfulness with a grouping winner selection mechanism.

A. Design Rational

If we want to provide exact truthfulness similar to *PADS-ADP*, we need to lose higher utility to achieve truthfulness. On the contrary, if the mechanism can tolerate a little loss of the truthfulness which is called ‘‘approximate truthfulness’’ as defined in Definition 3, we can design a more efficient mechanism with higher revenues and better privacy guarantees.

Without using the iterative method to choose the winners one by one, *PADS-DP* chooses the winners using a grouping method. The groups are calculated by the bid values which makes the mechanism individual rational for every user (See Definition 4). After that, *PADS-DP* chooses the winner group from the candidates using the exponential mechanism. In the next subsection, we will discuss the details of *PADS-DP*.

B. Design Detail

As shown in Figure 2, the mechanism needs to decide both the winners and the payments for the winners by the bids which is the input of the mechanism. Here, if we need to choose a payment scheme ensuring that every winner pays the same payment, $p_i = p, \forall x_i = 1$, then the possible outcome of the auction is not $\binom{n}{m}$ but n which is the number of the bids (each bid can be a possible outcome that $p = b_i$).

The basic idea of *PADS-DP* is to group the bids by the bid value. The mechanism is shown in Algorithm 2. First, we sort the bids in the descending order to build a set to denote the possible payment outcomes, $P = \{\rho_1, \rho_2, \dots, \rho_n\}$ and $\rho_1 \leq \rho_2 \leq \dots \leq \rho_n$. Then, we group the bids by the possible payments P . For each group, it has a payment scheme $p = \rho_i \in P$. Next, based on the payments, we select the winners from the highest bid until the lowest bid $b_j \geq \rho_i$ or until the number of winners is larger than or equal to m .

We use S_i to represent the candidates in the group. Therefore, in this condition, we have the following score function for each group:

$$F(S_i, \rho_i) = \rho_i |S_i| \quad (20)$$

For each group, based on the score function, we calculate the probability:

$$Pr_i = \frac{\exp(\frac{\epsilon}{2\Delta} \rho_i |S_i|)}{\sum_{\rho_j \in P} \exp(\frac{\epsilon}{2\Delta} \rho_j |S_j|)} \quad (21)$$

Finally, based on the probabilities, we randomly choose a group as the winners. The candidates S_i as the final winners and the payment is set to ρ_i . The computational complexity

Algorithm 2: PADS-DP Mechanism

Input : Type of the VM : k ;
of VMs: m ;
Buy bids: $\vec{b} = \{b_1, b_2, \dots\}$;
Output: Auction winners: S ;
Payment scheme: p ;

- 1 Initially, sort \vec{b} with descending order and generate $P = \vec{b}$;
- 2 **for** $\rho_i \in P$ **do**
- 3 **while** $b_j \geq \rho_i$ and s.t. Eq.(8) **do**
- 4 $S_i \leftarrow S_i \cup \{j\}$;
- 5 **end**
- 6 $Pr_i = \frac{\exp(\frac{\epsilon}{2\Delta} \rho_i |S_i|)}{\sum_{\rho_j \in P} \exp(\frac{\epsilon}{2\Delta} \rho_j |S_j|)}$;
- 7 **end**
- 8 Randomly select the winner group with probability \vec{Pr} ;
- 9 Assume the winner set is S_i ;
- 10 $S \leftarrow S_i$;
- 11 $p = \rho_i$;

of Algorithm 2 is $O(n * m)$ which is determined by the main loop from line 2 to 7 in Algorithm 2. It calculates the possible winners for each set S_i . The maximum number of winners is bounded by m and the number of the sets is n . So the computational complexity is $O(n * m)$.

After presenting the mechanism, we next provide the formal theoretical analysis of the desirable properties of *PADS-DP* mechanism. First, we prove that the PADS-DP mechanism is ϵ -differentially private in Theorem 7.

Theorem 7. The *PADS-DP* auction is ϵ -differentially private.

Proof. We denote \vec{b} and \vec{b}' as two bid profiles that differ in only one bidder’s bid. We use M to denote *PADS-DP* mechanism, $\forall p \in P$, we have:

$$\begin{aligned} & \frac{Pr[M(\vec{b}) = p]}{Pr[M(\vec{b}') = p]} \\ &= \frac{\exp(\frac{\epsilon}{2\Delta} p |S_i|) \sum_{\rho_j \in P} \exp(\frac{\epsilon}{2\Delta} \rho_j |S'_j|)}{\exp(\frac{\epsilon}{2\Delta} p |S'_i|) \sum_{\rho_j \in P} \exp(\frac{\epsilon}{2\Delta} \rho_j |S_j|)} \\ &\leq \exp(\frac{\epsilon}{2\Delta} p) \frac{\sum_{\rho_j \in P} \exp(\frac{\epsilon}{2\Delta} \rho_j (|S_j| + 1))}{\sum_{\rho_j \in P} \exp(\frac{\epsilon}{2\Delta} \rho_j |S_j|)} \quad (22) \\ &\leq \exp(\frac{\epsilon}{2}) \frac{\sum_{\rho_j \in P} \exp(\frac{\epsilon \rho_j |S_j| + \epsilon \Delta}{2\Delta})}{\sum_{\rho_j \in P} \exp(\frac{\epsilon}{2\Delta} \rho_j |S_j|)} \\ &= \exp(\frac{\epsilon}{2}) \exp(\frac{\epsilon}{2}) \\ &= \exp(\epsilon) \end{aligned}$$






Figure 3. Evaluation results for revenues






Figure 4. Evaluation results for social welfares

Therefore, we have:

$$Pr[M(\vec{b}) = p] \leq \exp(\epsilon) Pr[M(\vec{b}') = p], \forall p \in P \quad (23)$$

and we arrive at the conclusion that the *PADS-DP* mechanism is ϵ -differentially private. \square

Next, we prove that *PADS-DP* is $\epsilon\Delta$ -truthful.

Theorem 8. The *PADS-DP* auction is $\epsilon\Delta$ -truthful.

Proof. The step is similar to Theorem 7. We also use \vec{b} and \vec{b}' as two bid profiles that differ in only one bidder's bid. We use the conclusion Eq.23 which can be transformed to $Pr[M(\vec{b}) = p] \geq \exp(-\epsilon) Pr[M(\vec{b}') = p]$. Therefore, the expectation of any bidder i 's utility satisfies:

$$\begin{aligned} E_{p \sim M(\vec{b})}[u_i(p)] &= \sum_{p \in P} u_i(p) Pr[M(\vec{b}) = p] \\ &\geq \sum_{p \in P} u_i(p) \exp(-\epsilon) Pr[M(\vec{b}') = p] \\ &= \exp(-\epsilon) E_{p \sim M(\vec{b}')} [u_i(p)] \\ &\geq (1 - \epsilon) E_{p \sim M(\vec{b}')} [u_i(p)] \\ &= E_{p \sim M(\vec{b}')} [u_i(p)] - \epsilon E_{p \sim M(\vec{b}')} [u_i(p)] \end{aligned} \quad (24)$$

As the maximum utility of an individual user is bounded by Δ which is based on the utility function $u_i(p) = (v_i - p)x_i \leq \bar{b} - \underline{b} = \Delta$, we can get

$$E_{p \sim M(\vec{b})}[u_i(p)] \geq E_{p \sim M(\vec{b}')} [u_i(p)] - \epsilon\Delta \quad (25)$$

Therefore, with the Definition 3, we can conclude that *PADS-DP* is $\epsilon\Delta$ -truthful. \square

VII. EVALUATION

We have implemented *PADS* (both *PADS-ADP* and *PADS-DP*) in a simulator, and extensively evaluate their performance. On the CSP's side, the evaluation results show that *PADS-DP* can achieve relatively high revenues and social welfares compared with *PADS-ADP*. On the users' side, we analyze the payments that are paid by the users to obtain the resources. In addition, we measure the job completion rate showing that *PADS-DP* can get near optimum result while maintaining a relatively high privacy level.

A. Setup

The default setting of the experimental evaluation is described below: In our experiments, we assume that each CSP

Table I
DEFAULT CONFIGURATION

# of bidders	5000	ϵ	0.1
# of VMs	200	δ	0.25
bids	[0,1]	time slot length	5 minutes
simulate time	1 hour	job running time	10 minutes

provides one type of VMs to the users. The simulation time is set to one hour for the default setting and the time slot is set to be five minutes similar to the Amazon EC2 Spot Instance. We generate the bids and the jobs for each user to model the interactions between the users and the clients as shown in Figure 2. We assume that the jobs are batch processing jobs which can be interrupted during execution. The bids are generated from a uniform distribution $b_i \in [0, 1]$ and are equal to the true valuations of the VMs for the users as the mechanisms we evaluate in our experiments are all

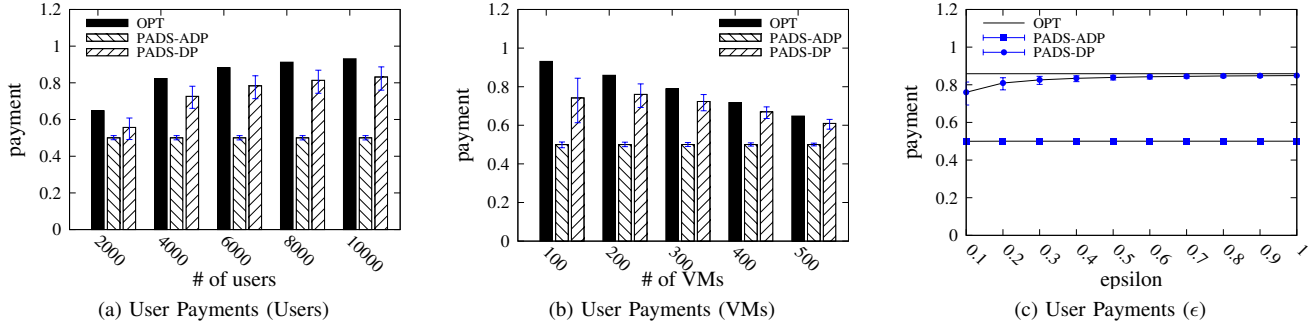


Figure 5. Evaluation results for user payments

truthful or approximately truthful mechanisms. The users do not change their bids during the one hour simulation period. The length of the jobs is set to be ten minutes, and therefore each job requires at least two time slots to complete. If a job is completed, the user client stops to bid for the VMs. All the experiment results are averaged for 100 trials. The error bar of each result represents the 95% confidence interval of the 100 trials.

B. Methodology

In our experiments, we compare *PADS-DP* and *PADS-ADP* with the VCG auction mechanism [17]–[19] (denoted as *OPT* in the results) which provides truthfulness but not differential privacy guarantees. The VCG mechanism is implemented in the simulator to satisfy the objectives and constraints in the dynamically priced resource allocation problem in clouds.

To evaluate the performance of *PADS*, we use the following four metrics:

- 1) *Revenue*: The revenue is calculated from the sum of all the payments from the users during the overall simulation time. We conduct three sets of experiments by increasing the number of users and the number of servers for different settings of ϵ .
- 2) *Social Welfare*: The social welfare is computed as the sum of the values of the users [27] which can be calculated as $\sum_i^n v_i x_i$. It is the basic metric to measure the economic efficiency of the auction mechanisms.
- 3) *User's Payment*: The user's payment is calculated as the average payment of each winner for each time slot.
- 4) *Completion Rate*: The completion rate of the jobs measures the fraction of the jobs that complete within their deadline.

C. Experiment Results

First, we measure the revenue earned by the CSPs in the auction process. We conduct three sets of experiments to evaluate the performance of *PADS-ADP* and *PADS-DP*. In our experiments, we study (i) the impact of the number of users, (ii) the influence of the number of the resources (VMs) on the auction performance and (iii) the impact of the ϵ parameter in the auction outcome.

From Figure 3a and 3b, we can see that *PADS-DP* achieves nearly the same revenue as *OPT*. In contrast, *PADS-ADP* attains only half the revenue of *OPT*. This observation reflects




Figure 6. Evaluation results for completion rate

the fact that *PADS-ADP* chooses the winners in a more random iterative manner compared to *PADS-DP* and hence its revenue is lower than *PADS-DP*. As shown in Figure 3a, we can also observe that when the number of bidders (users) is increased, the *OPT* and *PADS-DP* schemes can get higher revenue, but *PADS-ADP* maintains the same revenue. This is also due to the fact that *PADS-ADP* chooses winners more randomly compared to the other two mechanisms. In Figure 3b, we analyze the performance by increasing the number of VMs. We observe that all the three schemes namely *OPT*, *PADS-DP* and *PADS-ADP* obtain higher revenue as resources are increased. We study the influence of the ϵ parameter on the revenue in Figure 3c. We observe that for ϵ (from 0.1 to 1), the influence is perceptible for *PADS-DP*. However, it is insignificant for *PADS-ADP* as *PADS-ADP* uses an iterative process to achieve differential privacy, and in each iteration, the $\epsilon' = \frac{\epsilon}{(e-1)\Delta \ln(e/\delta)}$ is calculated based on ϵ . Since it is usually tenfold smaller than ϵ , it makes the scale of ϵ insensitive. For *PADS-DP*, with larger ϵ , the revenue approaches that of the *OPT* scheme with smaller variances (as shown by smaller error bars in the figure).

Next, we evaluate social welfare which is one of the metrics to measure the economic efficiency of the auction mechanisms [27]. In Figure 4a, we observe that *PADS-DP* achieves nearly similar social welfare as the number of users is increased. It is significantly better than that of *PADS-ADP* which maintains a constant social welfare irrespective of the number of users who participate in the auction. The social welfare measured in Figure 4b suggests that the schemes follow a similar trend as Figure 4a when the number of VMs is increased. From the results in Figure 4c for varying values of ϵ , we can observe and conclude that *PADS-DP* can achieve nearly the same social

welfare as that of *OPT* even for smaller values of ϵ .

In Figure 5, we measure the payment made by the users in the proposed schemes. We plot the average payment for each user in each time slot. We observe that in *PADS-DP*, users pay less than that of *OPT* (Figure 5a and 5b) for a different number of users and VMs. The trend in Figure 5c is also similar to that of Figure 3c.

Finally, in Figure 6, we consider the users' satisfaction which can be represented by the completion rate. In an auction-based resource allocation mechanism, if a user bids higher, the user should have a higher probability to obtain the resource. In Figure 6a and 6b, we plot the completion rate of each user marked as blue "X" and the results of *OPT* marked as solid lines. We can see that with higher ϵ , the completion rate of *PADS-DP* is closer to that of *OPT*. Since $\epsilon = 0.1$ is a significantly large privacy protection with differential privacy, the results demonstrate that *PADS-DP* can achieve relatively high differential privacy and user satisfaction simultaneously.

VIII. CONCLUSION

In this paper, we propose a strategy-proof differentially private auction mechanism for allocating dynamically priced resources in a cloud. We propose three approaches for the privacy-aware auction design problem using differential privacy based on exponential mechanism design. The first approach uses a straight-forward application of near optimum exponential mechanism that provides truthfulness and ϵ -differential privacy, but the mechanism is intractable for large-scale resource allocations. The second approach, *PADS-ADP*, uses an iterative algorithm to choose the winners of an auction and achieves (ϵ, δ) -differential privacy and runs in polynomial time. The third approach, *PADS-DP*, employs a grouping algorithm to generate the possible outcome groups and chooses the winner group using the exponential mechanism. We demonstrate that *PADS-DP* can achieve ϵ -differential privacy and $\epsilon\Delta$ -truthfulness. Experimental evaluation of the performance of the proposed mechanisms shows that the proposed techniques can guarantee differential privacy and truthfulness property of the auction while achieving closely similar performance in terms of revenue and social welfare as compared to traditional auctions.

REFERENCES

- [1] L. Columbus, "Roundup of cloud computing forecasts and market estimates q3 update, 2015," Sep. 2015. [Online]. Available: <http://www.forbes.com/sites/louiscolumbus/2015/09/27/roundup-of-cloud-computing-forecasts-and-market-estimates-q3-update-2015/#6a2e25b66c7a>
- [2] Amazon, "Amazon ec2 spot instances," Jan. 2017. [Online]. Available: <https://aws.amazon.com/ec2/spot/>
- [3] Amazon, "Amazon ec2 on-demand pricing," Jan. 2017. [Online]. Available: <https://aws.amazon.com/ec2/pricing/on-demand/>
- [4] L. Zheng, C. Joe-Wong, C. W. Tan, M. Chiang, and X. Wang, "How to bid the cloud," in *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*. ACM, 2015, pp. 71–84.
- [5] M. Yokoo, Y. Sakurai, and S. Matsubara, "The effect of false-name bids in combinatorial auctions: New fraud in internet auctions," *Games and Economic Behavior*, vol. 46, no. 1, pp. 174–188, 2004.
- [6] M. Naor, B. Pinkas, and R. Sumner, "Privacy preserving auctions and mechanism design," in *Proceedings of the 1st ACM conference on Electronic commerce*. ACM, 1999, pp. 129–139.
- [7] K. Q. Nguyen and J. Traoré, "An online public auction protocol protecting bidder privacy," in *Information Security and Privacy*. Springer, 2000, pp. 427–442.
- [8] H. Kikuchi, S. Hotta, K. Abe, and S. Nakanishi, "Distributed auction servers resolving winner and winning bid without revealing privacy of bids," in *Parallel and Distributed Systems: Workshops, Seventh International Conference on, 2000*. IEEE, 2000, pp. 307–312.
- [9] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*. IEEE, 2007, pp. 94–103.
- [10] J. McMillan, "Why auction the spectrum?" *Telecommunications policy*, vol. 19, no. 3, pp. 191–199, 1995.
- [11] Q. Huang, Y. Tao, and F. Wu, "Spring: A strategy-proof and privacy preserving spectrum auction mechanism," in *INFOCOM, 2013 Proceedings IEEE*. IEEE, 2013, pp. 827–835.
- [12] P. Cramton, "Spectrum auction design," *Review of Industrial Organization*, vol. 42, no. 2, pp. 161–190, 2013.
- [13] S. Liu, H. Zhu, R. Du, C. Chen, and X. Guan, "Location privacy preserving dynamic spectrum auction in cognitive radio network," in *Distributed Computing Systems (ICDCS), 2013 IEEE 33rd International Conference on*. IEEE, 2013, pp. 256–265.
- [14] F. Wu, Q. Huang, Y. Tao, and G. Chen, "Towards privacy preservation in strategy-proof spectrum auction mechanisms for noncooperative wireless networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 23, no. 4, pp. 1271–1285, 2015.
- [15] H. Huang, X.-Y. Li, Y.-e. Sun, H. Xu, and L. Huang, "Pps: Privacy-preserving strategyproof social-efficient spectrum auction mechanisms," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 26, no. 5, pp. 1393–1404, 2015.
- [16] C. Dwork, "Differential privacy: A survey of results," in *International Conference on Theory and Applications of Models of Computation*. Springer, 2008, pp. 1–19.
- [17] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of finance*, vol. 16, no. 1, pp. 8–37, 1961.
- [18] E. H. Clarke, "Multipart pricing of public goods," *Public choice*, vol. 11, no. 1, pp. 17–33, 1971.
- [19] T. Groves, "Incentives in teams," *Econometrica: Journal of the Econometric Society*, pp. 617–631, 1973.
- [20] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.
- [21] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography Conference*. Springer, 2006, pp. 265–284.
- [22] V. Rastogi and S. Nath, "Differentially private aggregation of distributed time-series with transformation and encryption," in *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*. ACM, 2010, pp. 735–746.
- [23] Z. Huang and S. Kannan, "The exponential mechanism for social welfare: Private, truthful, and nearly optimal," in *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*. IEEE, 2012, pp. 140–149.
- [24] A. Le Ny, "Introduction to (generalized) gibbs measures," *Ensaio Matemáticos*, vol. 15, pp. 1–126, 2008.
- [25] R. Zhu and K. G. Shin, "Differentially private and strategy-proof spectrum auction with approximate revenue maximization," in *Computer Communications (INFOCOM), 2015 IEEE Conference on*. IEEE, 2015, pp. 918–926.
- [26] R. Zhu, Z. Li, F. Wu, K. Shin, and G. Chen, "Differentially private spectrum auction with approximate revenue maximization," in *Proceedings of the 15th ACM international symposium on mobile ad hoc networking and computing*. ACM, 2014, pp. 185–194.
- [27] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic game theory*. Cambridge University Press Cambridge, 2007, vol. 1.
- [28] M. J. Osborne and A. Rubinstein, *A course in game theory*. MIT press, 1994.
- [29] A. Gupta, K. Ligett, F. McSherry, A. Roth, and K. Talwar, "Differentially private combinatorial optimization," in *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 2010, pp. 1106–1125.
- [30] A. Archer and É. Tardos, "Truthful mechanisms for one-parameter agents," in *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*. IEEE, 2001, pp. 482–491.