

DPDA: A Differentially Private Double Auction Scheme for Mobile Crowd Sensing

Wenqiang Jin*, Ming Li*, Linke Guo[†], Lei Yang*

* University of Nevada, Reno, U.S.

[†] Binghamton University, Binghamton, U.S.

Email: wenqjin@nevada.unr.edu, mingli@unr.edu, lguo@binghamton.edu, leiy@unr.edu

Abstract—Mobile crowd sensing (MCS) takes advantage of pervasive mobile devices that are equipped with multi-sensors to collect rich data of a certain geographic area. Because of the importance of incentivizing users to participate, auction-based open MCS markets have been proposed in past literature. Note that their focus is to achieve critical economic properties but fail to protect bid privacy. Although there are limited schemes dealing with this issue, they are designed only for single-side auctions and are unsuitable for double-side auctions whose properties are quite different. In this paper, inspired by uniform pricing and exponential mechanism, we propose a differentially private double auction (DPDA) scheme for MCS to protect bid privacy for both auction sides. In addition, the traditional economic properties, such as γ -truthfulness, individual rationality and budget balance, are guaranteed as well. Besides, we derive closed forms over the computation complexity and the approximate optimal platform revenue achieved by the scheme. Extensive simulations have been conducted on real-world datasets to validate the efficiency and effectiveness of DPDA.

I. INTRODUCTION

The exploding growth of mobile devices with diverse embed sensors (e.g., camera, gyroscope, 3D accelerometer and GPS) has triggered the thriving of human-centric mobile crowd sensing (MCS) which has been gradually recognized as a compelling paradigm for sharing and utilizing resources of personal hardwares to collect enormous data from surrounding environment. Typically, a MCS system is held by the cloud platform, accepting sensing tasks from requestors and collecting sensing data provided by mobile devices, namely sensing workers. Open markets like Crowdsignals [1] have emerged to facilitate interactions among requestors and workers for different task purposes.

Executing sensing tasks is resource-consuming for individual workers. Therefore, to elicit them to participate, auctions naturally serve as ideal means to provide suitable incentives to workers and to distribute sensing tasks among them. There have been some research efforts modeling MCS markets as reverse single-side auctions, where the platform plays as the auctioneer and workers are sellers to compete for sensing opportunities [2]–[6]. Realizing the competition also exists among task requestors, recent works [7]–[9] alternatively model MCS markets as double-side auctions, or called double auctions for short, where requestors play as buyers to purchase sensing efforts from workers. However, the main focus of these works is to achieve some economic properties, such as truthfulness and individual rationality, so as to guarantee robustness and sustainability of MCS markets. For instance,

in a truthful auction, a player's best strategy is to bid honestly; in the case of a double auction, a worker's (requestor's) ask (bid) is exactly his cost (value) toward a task.

Even with all these promising properties, as pointed out by [11], [12], an auction is vulnerable to inference attacks. A viable strategy for an adversary to perform such attacks is, for example, to try a different bid/ask at each auction round and then analyze the corresponding auction outcomes. In such a way, bids/asks from other victims can be revealed by differencing multiple rounds of outcomes. In addition, if the auction is truthful, disclosed bids/asks are exactly true costs/values from their owners. Leveraging these data, the adversary is able to infer his opponents' financial conditions, so as to bid strategically for beneficial gain. To address this issue in MCS, only a handful research [11], [12] has been conducted. The notion of differential privacy has been applied. However, they merely focus on tackling inference attacks in single-side auction scenarios, while how to resist these attacks in double-side markets in MCS has been neglected so far. Note that interactions among requestors, workers and the auctioneer in double auctions are much more complicated than single-side auctions, due to the involvement of requestors. Besides, in addition to determining winning workers and their payments, a double auction should also find out winning requestors and their charges. Moreover, additional economic properties, such as budget balance, should satisfy in double auctions. All these uniquenesses impose great challenges in designing double auction schemes for MCS that resist inference attacks.

In this paper, we develop a practical differentially private double auction (DPDA) scheme to combat inference attacks in MCS. Task requestors and workers participate in the market as buyers and sellers, respectively. In order to generate high revenue, the platform, who is the auctioneer, formulates a platform revenue maximization (PRM) problem. It takes into account various constraints from both sides and proves as NP-hard. To solve it efficiently, we propose a heuristic winner determination algorithm. A score function is carefully designed for matching requestors and their tasks to suitable workers so as to achieve the approximate optimal platform revenue. Then, uniform pricing is adopted to calculate payments to winning workers and charges to winning requestors, to guarantee economic properties, including truthfulness, individual rationality and budget balance. Moreover, as the most critical component, inspired by the exponential mechanism [13], we convert the winner selection and pricing into a probabilistic fashion for the effect that the change of one player's bid/ask only has limited

impact to an auction outcome. Hence, it thwarts the adversary from analyzing different outcomes for inference attacks. Our contributions in this paper are summarized as follows.

- We discuss how to achieve differential privacy for double-side markets in MCS, which has never been discussed before.
- Through rigid theoretical analysis, we prove that DPDA guarantees differential privacy, γ -truthfulness, individual rationality and budget balance. We also give closed forms over the computation complexity and approximate optimal platform revenue realized via our scheme.
- We extensively evaluate DPDA over real-world datasets.

The remainder of the paper is organized as follows. We review some relevant works and their deficiencies in Section II. Section III presents the system model and problem formulation. Design details of the proposed DPDA scheme are described in Section IV. Its theoretical analysis is provided in Section V, followed by evaluation in Section VI. Section VII concludes the entire paper.

II. RELATED WORK

Only a handful schemes have been proposed to protect bid privacy in MCS. Dimitriou and Krontiris [14] studied how to protect workers' bids from the platform. They leveraged some crypto primitives to guarantee bidders' anonymity and design a rewarding mechanism that enables winners to claim their reward without being linked to the data they contributed. However, this scheme does not discuss how to protect bid privacy from other bidders. Besides, inference attacks have also been ignored. Realizing this issue, Jin et al. [12] proposed a differentially private incentive mechanism that preserves the privacy of one's bid against other honest-but-curious workers, where they modeled the MCS market as the reverse combinatorial auction and targeted at minimizing the platform's total payment. In the same line of research, Lin et al. [11] also adopted the concept of differential privacy. Alternatively, they aimed to minimize the overall social cost. Note that all these works focus on single-side auctions in MCS, where the effect from the demand side, i.e., requestors, have not been taken into account. In this work, we consider a more practical yet complicated market for MCS; double-side auctions are applied to model the interactions among requestors, workers and the platform. Hence, how to resist inference attacks therein becomes a more challenging task that deserves a thorough investigation.

Bid privacy protection has also been studied in generic auctions. For example, Parkes et al. [15] presented a privacy-preserving and verifiable auction based on homomorphic cryptography. No party, including the auctioneer, receives any information about bids, and no bidder is able to change or repudiate any bid. Another privacy-preserving and verifiable auction scheme [16] is developed for online Ad exchanges. However, partial bid information is revealed once an auction closes in these schemes. Realizing that bids in previous auctions can be used in the future auctions, Nojournian and

Stinson [17] leveraged the verifiable secret sharing technique to make sure that bids, especially the losing ones, are kept private through the entire process of an auction. Note that the above schemes [15]–[17] do not consider bid privacy leakage caused by other players' inference attacks. McSherry et al. [18] are among the first to design differentially private auction mechanisms to defend against such attacks, followed by research [13], [19], [20]. However, these mechanisms are designed for generic auctions, which fail to take into account unique properties in MCS, such as sensing capabilities, preferences, travel budgets at workers and sensing quality requirements at requestors. With the involvement of these factors, the differentially private auction mechanism design for MCS will be significantly different.

It is also worth mentioning some other efforts on protecting worker privacy in MCS. For example, the work [21]–[23] target at hiding worker locations from the platform, while the work [24], [25] discuss how to preserve workers' identity and data privacy when submitting sensing results. Apparently, we are dealing with a totally different privacy concern from them.

III. SYSTEM MODEL AND PROBLEM STATEMENT

We consider a general open market for MCS which consists of a set of task requestors $\mathcal{R} = \{R_1, \dots, R_i, \dots, R_N\}$ who purchase sensing data for their tasks, a cloud platform who acts as the auctioneer, and a set of participatory workers $\mathcal{W} = \{W_1, \dots, W_j, \dots, W_M\}$ who compete for sensing opportunities. Typically, requestors have various sensing tasks, each with its own quality requirement and workers bear different sensing capabilities and preferences. Workers have to travel for a distance to perform tasks that are located in different geographical areas. In order to wisely assign requestors and their tasks to workers, a typical auction workflow can be summarized as follows.

- Each requestor $R_i \in \mathcal{R}$ submits his bid profile

$$F_i^R = \langle T_i, L_i, b_i \rangle$$

where T_i stands for the set of tasks to be completed at location L_i . b_i represents the maximal per-task payment he is willing to spend according to his own per-task valuation v_i for obtaining the sensing data. Following [26]–[28], we adopt normalized b_i and v_i , i.e., they take values from $(0, 1]$. Denote by F^R all requestors' bid profiles.

- Meanwhile, each worker $W_j \in \mathcal{W}$ submits his ask profile

$$F_j^W = \langle \tau_j, l_j, s_j \rangle$$

where τ_j indicates the set of tasks that he is willing to perform, l_j stands for his current location, and s_j represents the minimum per-task payment he accepts based on his per-task cost c_j . Similarly, s_j and c_j take values from $(0, 1]$. Denote by F^W all workers' ask profiles.

- Once collecting bid and ask profiles from players, the platform selects the winning requestors and workers, and determines the task assignment policy. Besides, it calculates a per-task payment p_j for each winning

worker W_j and a per-task charge a_i to each winning requestor R_i . Finally, the platform announces auction results.

A. Problem Formulation

The outcome of an auction heavily relies on its objective properties. In this work, as [7], [29], we intend to maximize the platform's revenue. Thus, the following platform revenue maximization (PRM) problem is formulated

$$\begin{aligned}
 \max \quad & \sum_{i=1}^N \sum_{j=1}^M (a_i - p_j) x_{i,j} \\
 \text{s.t.} \quad & \sum_{j=1}^M x_{i,j} \geq |T_i|, \quad \forall i \in [1, N] \\
 & \bigcup_{j: x_{i,j} \neq 0} \tau_j \supseteq T_i, \quad \forall i \in [1, N] \\
 & \rho(r_j) \leq D_j, \quad \forall j \in [1, M] \\
 & x_{i,j} \in \mathbb{Z}, \quad a_i, p_j \in \Gamma, \quad \forall i \in [1, N], \forall j \in [1, M].
 \end{aligned} \tag{1}$$

Denote by r_j the set of tasks that the platform pair to worker W_j . $\rho(r_j)$ represents the shortest accumulated path for completing all tasks in r_j . Also, it is practical to assume that each worker W_j has his travel budget D_j . The decision variables of PRM include $x_{i,j}$'s, a_i 's and p_j 's, among which $x_{i,j}$ is an integer, denoting the number of tasks W_j executes for R_i . In the above problem, constraints (1) and (2) require that a winning requestor's sensing tasks should be fully completed. Constraint (3) states that a worker's accumulated traveling distance should not surpass its travel budget. Besides, we define a finite price set Γ that contains all feasible values of payments p_j 's and charges a_i 's.

B. Objectives and Assumptions

The primary goal of this work is to defend against inference attacks in MCS auction markets. As mentioned above, in an inference attack, one player (i.e., a requestor or worker) can infer others' bids/asks by changing its own bid/ask in multiple auction rounds and analyzing the corresponding auction outputs. With others' bid/ask information, one can use it as auxiliary knowledge to further infer opponents' financial conditions and business strategies so as to play with favorable actions for beneficial gain. Therefore, we aim to design a differentially private mechanism to address this issue. The formal definition of differential privacy is given below.

Definition 1. Differential Privacy (revised from [18], [30]). We denote the proposed auction mechanism as a function $M(\cdot)$ that maps input bids ($\mathbf{b}$) and asks ($\mathbf{s}$) to charges ($\mathbf{a}$) and payments ($\mathbf{p}$). Then $M(\cdot)$ guarantees ϵ -differential privacy for its inputs, if and only if for any possible set of payments and charges $\{\mathbf{a}, \mathbf{p}\}$, and any two bid and ask sets $\{\mathbf{b}, \mathbf{s}\}$ and $\{\mathbf{b}', \mathbf{s}'\}$ (or $\{\mathbf{b}, \mathbf{s}'\}$) that only differ in one element, we have

$$\Pr[M(\{\mathbf{b}, \mathbf{s}\}) \in \{\mathbf{a}, \mathbf{p}\}] \leq \exp(\epsilon) \Pr[M(\{\mathbf{b}', \mathbf{s}'\}) \in \{\mathbf{a}, \mathbf{p}\}]$$

or

$$\Pr[M(\{\mathbf{b}, \mathbf{s}\}) \in \{\mathbf{a}, \mathbf{p}\}] \leq \exp(\epsilon) \Pr[M(\{\mathbf{b}, \mathbf{s}'\}) \in \{\mathbf{a}, \mathbf{p}\}]$$

where ϵ is a constant usually referred as privacy budget.

Definition 1 means that changing a player's bid/ask does not result in significant changes of the final payment and charge profile. Thus, any player alone cannot successfully launch an inference attack through manipulating his bid/ask and analyzing the differences of auction outputs.

In this work, we aim to protect one's bid privacy from other players who are semi-honest. Namely, they behave honestly but are curious in acquiring others' bid information. Besides, like existing differential privacy mechanism design [11], [12], we assume that the central controller, i.e., platform here, is trustworthy; it does not leverage its knowledge over players' inputs and manipulate auction outcomes. Thus, we do not intend to protect bid privacy from it.

In addition to protecting players' bid privacy, we don't want to sacrifice some critical economic properties, such as truthfulness, individual rationality and budget balance. We first present the formal definition of a player's utility.

Definition 2. A Player's Utility. A requestor R_i 's utility u_i^R ($R_i \in \mathcal{R}$) is the difference between his valuation and charge toward tasks,

$$u_i^R = \begin{cases} \sum_{j \in [1, M]} x_{i,j} (v_i - a_i), & \text{if } R_i \text{ is a winner} \\ 0, & \text{Otherwise} \end{cases}$$

Also, a worker W_j 's utility u_j^W ($W_j \in \mathcal{W}$) is the difference between his payment and cost toward tasks,

$$u_j^W = \begin{cases} \sum_{i \in [1, N]} x_{i,j} (p_j - c_j), & \text{if } W_j \text{ is a winner} \\ 0, & \text{Otherwise} \end{cases}$$

Typically, auction players are strategic in a sense that they manipulate their bids/asks to win an auction. Truthfulness ensures players bid strictly following their values/costs. Under the framework of differential privacy, following [31], we adopt a notion of relaxed truthfulness. It guarantees no player can gain more than γ utility by bidding untruthfully.

Definition 3. γ -Truthfulness. An auction is γ -truthfulness for requestors, if and only if, for each requestor $R_i \in \mathcal{R}$, we have

$$E[u_i^R(v_i, \mathbf{b}_{-i})] \geq E[u_i^R(b_i, \mathbf{b}_{-i})] - \gamma,$$

where $b_i \neq v_i$ and $\mathbf{b}_{-i}$ denotes other requestors' bids. And for each worker $W_j \in \mathcal{W}$, we have

$$E[u_j^W(c_j, \mathbf{s}_{-j})] \geq E[u_j^W(s_j, \mathbf{s}_{-j})] - \gamma,$$

where $s_i \neq c_i$ and $\mathbf{s}_{-j}$ denotes other workers' asks.

Moreover, the utility received by an individual player should be nonnegative, which is defined as individual rationality.

Definition 4. Individual Rationality. Our auction mechanism is individual rationality, if and only if, for each requestor $R_i \in \mathcal{R}$ and worker $W_j \in \mathcal{W}$, we have

$$u_i^R \geq 0 \quad \text{and} \quad u_j^W \geq 0.$$

The platform revenue should be nonnegative to ensure the market sustainability, which is referred as budget balance.

Definition 5. Budget Balance. An auction is budget-balanced, if and only if, the auctioneer's revenue is nonnegative, i.e.,

$$\sum_{i=1}^N \sum_{j=1}^M (a_i - p_j) x_{i,j} \geq 0.$$

For time-sensitive MCS, it is critical to solve PRM efficiently. Unfortunately, it turns out to be NP-hard.

Theorem 1. The PRM problem is NP-hard.

Proof: For the analysis, we first degenerate the problem to a special case in which there is only one worker in the system and it can perform all requestors' tasks. Under this setting, the PRM problem is transferred into an *orienteering problem*, in which tasks can be modeled as vertices V of a graph $G(V, E)$. Each vertex is assigned with a reward, which is the revenue for completing the corresponding task. There is an edge connecting arbitrary two vertices of V . The weight for each edge is directly the distance between two task locations. For the *orienteering problem*, it tries to find a path in the graph that originates at the beginning vertex with the total length no larger than a certain threshold, such that the accumulated reward for arriving each vertex along the path is maximized. In our problem it is equivalent to identify the worker's best task set and the travel path such that the platform revenue is maximized. According to [32], [33], an *orienteering problem* is NP-hard. Since a simplified version of PRM is already NP-hard, the PRM is NP-hard as well. ■

Theorem 1 tells that it is time consuming to optimally solve PRM. Therefore, it is desirable to design a heuristic algorithm to solve it efficiently.

In conclusion, our design goal in this work is to achieve differential privacy, γ -truthfulness and individual rationality, and budget balance in MCS auctions. Besides, we also aim to derive approximate optimal platform revenue with computation efficiency via the heuristic algorithm.

IV. DPDA SCHEME DESIGN

Our proposed scheme consists of two procedures, winner determination and pricing. In the following, we elaborate them with details.

A. Winner Determination

Recall that Γ stands for the finite set including all possible charges from requestors and payments to workers. We first present the following definition which is critical in our scheme design.

Definition 6. Given the pricing pair $(a, p) \in \Gamma^2$ ($a \geq p$), $\mathcal{R}^{(a)} \subset \mathcal{R}$ is a set of requestors whose bids are no smaller than a ; $\mathcal{W}^{(p)} \subset \mathcal{W}$ is a set of workers whose asks are no larger than p .

Once receiving bid/ask profiles from all players, the platform derives $\mathcal{R}^{(a)}$ and $\mathcal{W}^{(p)}$ for each pair $(a, p) \in \Gamma^2$. We

Algorithm 1 Winner Determination

Input: $\Gamma^2, \mathcal{R}, \mathcal{W}, F^R, F^W$

Output: $\mathcal{R}^{(a)}, \mathcal{W}^{(p)}, \mathbf{x}_p^a, \Delta_p^a, \forall (a, p) \in \Gamma^2$

```

1: for all  $(a, p) \in \Gamma^2$  do
2:   Derive  $\mathcal{R}^{(a)}$  and  $\mathcal{W}^{(p)}$ ;
3:   Sort  $\mathcal{R}^{(a)}$  according to  $|T_i|$ ;
4:   for all  $R_i \in \mathcal{R}^{(a)}$  do
5:     while  $T_i \neq \emptyset$  and  $\bigcup_{j: W_j \in \mathcal{W}^{(p)}} \tau_j \geq T_i$  and
            $\max_{j: W_j \in \mathcal{W}^{(p)}} \{D_j - d_{i,j}\} \geq 0$  do
6:       for all  $W_j \in \mathcal{W}^{(p)}$  do
7:         Calculate  $score(i, j)$  following (4);
8:       end for
9:        $s = \arg \min_{j: score(i, j) \geq 0} score(i, j)$ ;
10:       $x_{i,s} \leftarrow x_{i,s} + |T_i \cap \tau_s|$ ,  $T_i \leftarrow T_i \setminus T_i \cap \tau_s$ ;
11:       $D_s \leftarrow D_s - d_{i,s}$ ,  $l_s \leftarrow l_i$ ;
12:       $\mathcal{R}^{(a)} \leftarrow \mathcal{R}^{(a)} \cup R_i$ ,  $\mathcal{W}^{(p)} \leftarrow \mathcal{W}^{(p)} \cup W_s$ ;
13:    end while
14:  end for
15:   $\Delta_p^a = \sum_{i: R_i \in \mathcal{R}^{(a)}} |T_i|$ ,  $\mathbf{x}_p^a = \{x_{i,j}\}$ ;
16: end for
```

now discuss how to determine the corresponding winning requestors $\mathcal{R}^{(a)} \subset \mathcal{R}^{(a)}$ and workers $\mathcal{W}^{(p)} \subset \mathcal{W}^{(p)}$. In order to admit more tasks so as to potentially bring in higher revenue for the platform, our heuristic algorithm tends to select winning requestors that have more tasks to execute; among $\mathcal{R}^{(a)}$, the requestor with the largest amount of tasks is processed in the first round, then the one with the second largest amount of tasks is processed next, etc. This iteration continues until all requestors in $\mathcal{R}^{(a)}$ have been examined. For each selected requestor $R_i \in \mathcal{R}^{(a)}$, the next step is to assign workers to execute its tasks. For this purpose, we design a function $score(i, j)$ that evaluates the fitness of assigning any worker W_j to requestor R_i

$$score(i, j) = \frac{|T_i| - |T_i \cap \tau_j|}{D_j - d_{i,j}}, \quad (4)$$

where $d_{i,j}$ represents the distance for W_j to travel to execute R_i 's tasks. Generally, the smaller value $score(i, j)$ is, the more tasks W_j can accomplish for R_i within a shorter travel distance. Besides, we set $d_{i,j} = \infty$ if $T_i \cap \tau_j = \emptyset$. Hence, if $score(i, j)$ is calculated negative, it indicates that W_j cannot accomplish R_i 's tasks within his travel budget D_j , or W_j is unwilling to execute R_i 's tasks. R_i 's tasks are assigned in an iterative manner. In each iteration, a winning worker W_s is selected in a way that its non-negative $score(i, s)$ is the minimum among all workers. Thereafter, the corresponding parameters are updated, including W_s 's task assignment $x_{i,s}$, remaining travel budget D_s , its current location l_s , R_i 's remaining task set T_i , and winning worker set $\mathcal{W}^{(b)}$. This worker assignment process for R_i continues until all its tasks have been distributed, or they cannot be satisfied. For the former case, R_i is a winning requestor belonging to $\mathcal{R}^{(a)}$; otherwise, it loses. We further denote by Δ_p^a the total number

of admitted tasks from all winning requestors. It is calculated as $\Delta_p^a = \sum_{i: R_i \in \bar{\mathcal{R}}^{(a)}} |T_i|$. Δ_p^a plays a critical role in the pricing procedure, which will be clear soon. Besides, we denote by $\mathbf{x}_p^a$ the requestor-worker assignment vector under (a, p) , $\mathbf{x}_p^a = \{x_{i,j} | i \in [1, N], j \in [1, M]\}$. Up to now, the winner set $\bar{\mathcal{R}}^{(a)}$ and $\bar{\mathcal{W}}^{(p)}$, and the requestor-worker assignment policy $\mathbf{x}_p^a$ have been determined for (a, p) . The same computation is conducted for the rest $(a', p') \in \Gamma^2$. Algorithm 1 gives an overview of our winner determination procedure.

From the description above, this procedure identifies the potential winners and their assignment policy under each pair $(a, p) \in \Gamma^2$. The remaining task is to determine which pair of (a, p) to adopt. As once it is fixed, so do the winners. As this part is related to the pricing procedure, we leave its discussion therein.

B. Pricing

Our pricing scheme is motivated by both the *uniform pricing* in auctions [34] and the *exponential mechanism* [18] in achieving differential privacy. For a (a, p) , each winning requestor in $\bar{\mathcal{R}}^{(a)}$ pays a uniform price a per task; each winning worker in $\bar{\mathcal{W}}^{(p)}$ gets paid with a uniform price p per task. To determine which (a, p) serves as the final pricing pair, we calculate the selection probability

$$\Pr[(a, p)] = \frac{\exp[\frac{\epsilon}{2K}(a-p)\Delta_p^a]}{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}^{\tilde{a}}]}, \quad (5)$$

where $K = \sum_{i: R_i \in \mathcal{R}} |T_i|$, i.e., the total number of tasks in the market. Recall that ϵ is the privacy budget we aim to achieve. Apparently, this selection probability is proportional to the exponential value $\exp[\frac{\epsilon}{2K}(a-p)\Delta_p^a]$. It can be interpreted in this way: (a, p) is more likely to be selected if winning requestors pay more and winning workers get paid less per task (i.e., a larger $(a-p)$), and it results in larger amount of tasks to be executed (i.e., a larger Δ_p^a).

Finally, once (a, p) is fixed following the probability (5), the winners $\bar{\mathcal{R}}^{(a)}$ and $\bar{\mathcal{W}}^{(p)}$, and their assignment policy $\mathbf{x}_p^a$ will be fixed as well. For each winning requestor $R_i \in \bar{\mathcal{R}}^{(a)}$, it pays $\sum_{j \in [1, M]} a \cdot x_{i,j}$ in total; for each winning worker $W_j \in \bar{\mathcal{W}}^{(p)}$, it gets paid at $\sum_{i \in [1, N]} p \cdot x_{i,j}$ in total.

V. ANALYSIS

In this section, we provide formal theoretical analysis on desirable properties of our proposed DPDA scheme.

Theorem 2. *The DPDA scheme is ϵ -differentially private.*

Proof: This property should be examined for both requestors and workers.

For requestors, denote by $\mathbf{b}$ and $\mathbf{b}'$ two bid sets from requestors that only differ in one requestor's bid. Besides, let

$\Delta_p'^a$ be the number of assigned tasks under $\mathbf{b}'$ for a given pricing pair $(a, p) \in \Gamma^2$. We have

$$\begin{aligned} & \frac{\Pr[M(\mathbf{b}, \mathbf{s}) = (a, p)]}{\Pr[M(\mathbf{b}', \mathbf{s}) = (a, p)]} \\ &= \frac{\exp[\frac{\epsilon}{2K}(a-p)\Delta_p^a]}{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}^{\tilde{a}}]} / \frac{\exp[\frac{\epsilon}{2K}(a-p)\Delta_p'^a]}{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}'^a]} \\ &= \exp[\frac{\epsilon(\Delta_p^a - \Delta_p'^a)}{2K}(a-p)] \cdot \frac{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}'^a]}{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}^{\tilde{a}}]} \\ &\stackrel{\textcircled{1}}{\leq} \exp[\frac{\epsilon}{2}(a-p)] \frac{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})(\Delta_p^a + K)]}{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}^{\tilde{a}}]} \\ &= \exp[\frac{\epsilon}{2}(a-p)] \frac{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}^{\tilde{a}}] \cdot \exp[\frac{\epsilon}{2}(\tilde{a}-\tilde{p})]}{\sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}^{\tilde{a}}]} \\ &\stackrel{\textcircled{2}}{\leq} \exp[\frac{\epsilon}{2}(a-p)] \cdot \exp(\frac{\epsilon}{2}) \leq \exp(\frac{\epsilon}{2}) \cdot \exp(\frac{\epsilon}{2}) = \exp(\epsilon) \end{aligned}$$

where $\textcircled{1}$ comes from $\Delta_p^a - \Delta_p'^a \leq K$, and $\textcircled{2}$ is due to $\exp[\frac{\epsilon}{2}(\tilde{a}-\tilde{p})] \leq \exp(\frac{\epsilon}{2})$ as $\tilde{a}-\tilde{p} \leq 1$. From the above expression, we have

$$\Pr[M(\mathbf{b}, \mathbf{s}) = (a, p)] \leq \exp(\epsilon) \cdot \Pr[M(\mathbf{b}', \mathbf{s}) = (a, p)].$$

According to Definition 1, we can infer that DPDA is ϵ -differentially private to requestors.

Following the similar idea above, denote by $\mathbf{s}$ and $\mathbf{s}'$ two ask sets from workers that only differ in one worker's ask. We have

$$\Pr[M(\mathbf{b}, \mathbf{s}) = (a, p)] \leq \exp(\epsilon) \cdot \Pr[M(\mathbf{b}, \mathbf{s}') = (a, p)],$$

i.e., DPDA is ϵ -differentially private to workers.

To sum up, DPDA is ϵ -differentially private to both requestors and workers. $\blacksquare$

To evaluate the γ -truthfulness, we first give following lemmas.

Lemma 1. *For each requestor $R_i \in \mathcal{R}$, given a pricing pair $(a, p) \in \Gamma^2$, we always have $u_i^R(v_i, \mathbf{b}_{-i}) \geq u_i^R(b_i, \mathbf{b}_{-i})$ via our DPDA scheme.*

Proof: We first consider the scenario where $v_i > b_i$.

- **Case 1:** R_i wins with both v_i and b_i . First of all, we can infer that $\bar{\mathcal{R}}^{(a)}(v_i, \mathbf{b}_{-i}) = \bar{\mathcal{R}}^{(a)}(b_i, \mathbf{b}_{-i})$ according to the winner determination rule (Algorithm 1), and thus $\mathbf{x}_p^a(v_i, \mathbf{b}_{-i}) = \mathbf{x}_p^a(b_i, \mathbf{b}_{-i})$. Hence, we have

$$u_i^R(v_i, \mathbf{b}_{-i}) = \sum_{j: W_j \in \bar{\mathcal{W}}} x_{i,j}(v_i - a) = u_i^R(b_i, \mathbf{b}_{-i}).$$

- **Case 2:** R_i wins with v_i but loses with b_i . Therefore,

$$u_i^R(v_i, \mathbf{b}_{-i}) > u_i^R(b_i, \mathbf{b}_{-i}) = 0.$$

- **Case 3:** R_i loses with v_i but wins with b_i . There are two sub-cases to consider. Sub-case I: $R_i(v_i), R_i(b_i) \in \overline{\mathcal{R}}^{(a)}$. We have $x_p^a(v_i, \mathbf{b}_{-i}) = x_p^a(b_i, \mathbf{b}_{-i})$ according to Algorithm 1. Hence, it is impossible to have “ R_i loses with v_i but wins with b_i ” under this sub-case. Sub-case II: $R_i(v_i) \notin \overline{\mathcal{R}}^{(a)}$ and $R_i(b_i) \in \overline{\mathcal{R}}^{(a)}$. According to how $\overline{\mathcal{R}}^{(a)}$ is formed, we have $v_i < a \leq b_i$, which contradicts with the scenario $v_i > b_i$. Based on the discussion over these two sub-cases, we can infer that the statement of Case 3 does not exist.
- **Case 4:** R_i loses with both v_i and b_i . Hence,

$$u_i^R(v_i, \mathbf{b}_{-i}) = u_i^R(b_i, \mathbf{b}_{-i}) = 0.$$

From the discussion above, we conclude that $u_i^R(v_i, \mathbf{b}_{-i}) \geq u_i^R(b_i, \mathbf{b}_{-i})$ when $v_i > b_i$. The proof is similar when $v_i < b_i$, which is omitted due to space limit.

To sum up, $u_i^R(v_i, \mathbf{b}_{-i}) \geq u_i^R(b_i, \mathbf{b}_{-i})$ always holds. ■

Lemma 2. For any worker $W_i \in \mathcal{W}$, given a pricing pair $(a, p) \in \Gamma^2$, we always have $u_i^W(c_i, \mathbf{s}_{-i}) \geq u_i^W(s_i, \mathbf{s}_{-i})$ via our DPDA scheme.

Proof: The proof is very similar to that for Lemma 1. Due to the space limit, we omit its discussion here. ■

Theorem 3. The DPDA scheme is γ -truthful.

Proof: Denote by $\mathbf{b}$ and $\mathbf{b}'$ two bid sets from requestors that only differ in one requestor's bid. For each requestor $R_i \in \mathcal{R}$, we have

$$\begin{aligned} E[u_i^R(v_i, \mathbf{b}_{-i})] &= \sum_{(a,p) \in \Gamma^2} u_i^R(v_i, \mathbf{b}_{-i}) \Pr[M(\mathbf{b}, \mathbf{s}) = (a, p)] \\ &\stackrel{\textcircled{1}}{\geq} \exp(\epsilon) \sum_{(a,p) \in \Gamma^2} u_i^R(b_i, \mathbf{b}_{-i}) \Pr[M(\mathbf{b}', \mathbf{s}) = (a, p)] \\ &\geq \exp(-\epsilon) \sum_{(a,p) \in \Gamma^2} u_i^R(b_i, \mathbf{b}_{-i}) \Pr[M(\mathbf{b}', \mathbf{s}) = (a, p)] \\ &= \exp(-\epsilon) E[u_i^R(b_i, \mathbf{b}_{-i})] \stackrel{\textcircled{2}}{\geq} (1 - \epsilon) E[u_i^R(b_i, \mathbf{b}_{-i})] \\ &\stackrel{\textcircled{3}}{\geq} E[u_i^R(b_i, \mathbf{b}_{-i})] - \epsilon \cdot \max_{i \in [1, N]} \{|T_i|\} \end{aligned}$$

where $\textcircled{1}$ can be simply derived from Theorem 2 and Lemma 1. $\textcircled{2}$ is because $\exp(-\epsilon) \geq 1 - \epsilon$. Note that the maximal possible utility for R_i should be no larger than $\max_{i \in [1, N]} \{|T_i|\} \cdot (v_i - a)$. Besides, $\max_{i \in [1, N]} \{|T_i|\} \cdot (v_i - a) \leq \max_{i \in [1, N]} \{|T_i|\}$ as $v_i, a \in [0, 1)$. Thus, $\textcircled{3}$ holds. From the discussion above, we conclude that requestors are $\epsilon \cdot \max_{i \in [1, N]} \{|T_i|\}$ -truthful in DPDA.

Similarly, based on Theorem 2 and Lemma 2, for each worker $W_j \in \mathcal{W}$

$$E[u_j^W(c_j, \mathbf{s}_{-j})] \geq E[u_j^W(s_j, \mathbf{s}_{-j})] - \epsilon \cdot \max_{j \in [1, M]} \{\tau_j\}.$$

Hence, workers are $\epsilon \cdot \max_{j \in [1, M]} \{\tau_j\}$ -truthful in DPDA.

By setting $\gamma = \max\{\epsilon \cdot \max_{i \in [1, N]} \{|T_i|\}, \epsilon \cdot \max_{j \in [1, M]} \{\tau_j\}\}$, DPDA guarantees γ -truthfulness for both requestors and workers. ■

Theorem 3 states that no player can gain in his expected utility with γ by bidding untruthfully. Therefore, it is reasonably to consider that they bid truthfully in our DPDA scheme, i.e., $b_i = v_i$ ($i \in [1, N]$), $s_j = c_j$ ($j \in [1, M]$).

Theorem 4. The DPDA scheme is individual rational.

Proof: Regarding $R_i \in \mathcal{R}$, for a given pricing pair (a, p) , its utility is calculated as

$$u_i^R = \sum_{j \in [1, M]} x_{i,j}(v_i - a) \stackrel{\textcircled{1}}{=} \sum_{j \in [1, M]} x_{i,j}(b_i - a) \stackrel{\textcircled{2}}{\geq} 0$$

where $\textcircled{1}$ is due to Theorem 3 and $\textcircled{2}$ can be directly derived from our winner determination rule (Algorithm 1).

Similarly, regarding $W_i \in \mathcal{W}$, for a given pricing pair (a, p) , is easy to deduce

$$u_j^W = \sum_{i: R_i \in R} x_{i,j}(p_j - c_j) \geq 0.$$

According to Definition 4, the DPDA scheme is individual rational. ■

Theorem 5. The DPDA scheme is budget-balanced.

Proof: The platform revenue produced by DPDA can be calculated as

$$\sum_{i=1}^N \sum_{j=1}^M (a - p) x_{i,j} \geq 0.$$

The inequality directly comes from the condition $a \geq p$. According to Definition 5, the DPDA scheme is budget-balanced. ■

Recall that another design goal is to achieve computation efficiency, as the original PRM problem is NP-hard. The following theorem states that our DPDA scheme has polynomial-time computation complexity that is related to the requestor number N , worker number M , total task number K , and the cardinality of the pricing set Γ .

Theorem 6. The computation complexity of our DPDA scheme is upper bounded by $\mathcal{O}((\max\{N \log N, KM\}) \cdot |\Gamma^2|)$.

Proof: The computation complexity of DPDA is dominated by the winner determination process (Algorithm 1), whose main loop includes $|\Gamma^2|$ iterations. For each iteration, the computation complexity for sorting $\mathcal{R}^{(a)}$ (line 3) is $\mathcal{O}(N \log N)$ in general. From line 4 to line 14, the algorithm assigns workers to tasks from each requestor in $\mathcal{R}^{(a)}$. Specifically, it involves M comparisons in line 9 with at most K iterations for assigning all K tasks to corresponding workers. Its computation cost is thus upper bounded by $\mathcal{O}(KM)$. Therefore, the computation complexity of DPDA is upper bounded by $\mathcal{O}((\max\{N \log N, KM\}) \cdot |\Gamma^2|)$. ■

In the following, we analyze the approximation ratio of the platform revenue (denoted by PR), generated by the

DPDA scheme, to the optimal platforms revenue (denoted by OPT), generated by directly solving PRM optimally. This ratio indicates how much revenue our scheme sacrifices in order to achieve all objective properties. To analyze this ratio, we first give the following lemma.

Lemma 3. Define OPT^* , the maximal revenue the platform can obtain when adopting arbitrary uniform pricing $(a, p) \in \Gamma^2$. Then

$$\frac{\max_{(a,p) \in \Gamma^2} \Delta_p^a(a-p)}{K} OPT \leq OPT^* \leq OPT.$$

Proof: First of all, we know OPT^* serves as a lower bound of OPT .

Recall that K stands for the total number of tasks in the market. We have

$$OPT \stackrel{\textcircled{1}}{\leq} K \cdot \max_{(\tilde{a}, \tilde{p}) \in \Gamma^2} (\tilde{a} - \tilde{p}) \stackrel{\textcircled{2}}{\leq} K \quad (6)$$

where $\textcircled{1}$ can be simply derived from the formulation of PRM and $\textcircled{2}$ is due to $\tilde{a}, \tilde{p} \in (0, 1]$. Besides, from the definition of OPT^* , we have $OPT^* \geq \max_{(a,p) \in \Gamma^2} \Delta_p^a(a-p)$. Together with (6), it derives

$$OPT^* \cdot 1 \geq \max_{(a,p) \in \Gamma^2} \Delta_p^a(a-p) \cdot \frac{OPT}{K}$$

which ends the proof. ■

Lemma 3 not only serves as a critical step for deriving Theorem 7, but also tells how much platform's revenue the uniform pricing trades for achieving truthfulness (regardless of differential privacy). Generally speaking, when an auction adopts uniform pricing, it can achieve the truthfulness property. This statement can be examined via the proof of Lemma 1 and Theorem 3 for our DPDA scheme which is exactly developed based on uniform pricing.

Theorem 7. Let $E[PR]$ be the expected platform's revenue achieved via our DPDA scheme. Then the relation between $E[PR]$ and OPT is given by

$$E[PR] \geq \frac{\max_{(a,p) \in \Gamma^2} \Delta_p^a(a-p)}{K} OPT - \frac{6K}{\epsilon} \ln(e + \frac{\epsilon OPT |\Gamma^2|}{2K}).$$

Proof: For some constant $t > 0$, define $S_t = \{(a, p) : PR > OPT^* - t\}$ and $\bar{S}_{2t} = \{(a, p) : PR < OPT^* - 2t\}$. We have

$$\begin{aligned} & \frac{\Pr(\bar{S}_{2t})}{\Pr(S_t)} \\ &= \frac{\sum_{(a,p) \in \bar{S}_{2t}} \exp[\frac{\epsilon}{2K}(a-p)\Delta_p^a] / \sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}^{\tilde{a}}]}{\sum_{(a,p) \in S_t} \exp[\frac{\epsilon}{2K}(a-p)\Delta_p^a] / \sum_{(\tilde{a}, \tilde{p}) \in \Gamma^2} \exp[\frac{\epsilon}{2K}(\tilde{a}-\tilde{p})\Delta_{\tilde{p}}^{\tilde{a}}]} \\ &= \frac{\sum_{(a,p) \in \bar{S}_{2t}} \exp[\frac{\epsilon}{2K}(a-p)\Delta_p^a]}{\sum_{(a,p) \in S_t} \exp[\frac{\epsilon}{2K}(a-p)\Delta_p^a]} \leq \frac{|\bar{S}_{2t}| \exp[\frac{\epsilon}{2K}(OPT^* - 2t)]}{|S_t| \exp[\frac{\epsilon}{2K}(OPT^* - t)]} \\ &= \exp(-\frac{\epsilon t}{2K}) \frac{|\bar{S}_{2t}|}{|S_t|} \leq \exp(-\frac{\epsilon t}{2K}) \frac{|\Gamma^2|}{|S_t|}. \end{aligned}$$

Together with the fact $\Pr(S_t) \leq 1$, it derives $\Pr(S_{2t}) = 1 - \Pr(\bar{S}_{2t}) \geq 1 - \exp(-\frac{\epsilon t}{2K}) \frac{|\Gamma^2|}{|S_t|}$. It implies that DPDA selects the pricing pair (a, p) such that the corresponding PR achieves at least $OPT^* - 2t$ with a probability at least $1 - \exp(-\frac{\epsilon t}{2K}) \frac{|\Gamma^2|}{|S_t|}$.

In addition, if t satisfies $t \geq \frac{2K}{\epsilon} \ln \frac{|\Gamma^2| OPT^*}{t |S_t|}$, then $\Pr(S_{2t}) \geq 1 - \frac{t}{OPT^*}$, and

$$\begin{aligned} E[PR] &\geq \sum_{(a,p) \in S_{2t}} (OPT^* - 2t) \Pr[M(\mathbf{b}, \mathbf{s}) = (a, p)] \\ &\geq (OPT^* - 2t)(1 - \frac{t}{OPT^*}) = OPT^* - 3t + \frac{t^2}{OPT^*} \\ &\geq OPT^* - 3t. \end{aligned} \quad (7)$$

Setting $t = \frac{2K}{\epsilon} \ln(e + \frac{\epsilon OPT^* |\Gamma^2|}{2K})$, we have

$$\begin{aligned} & \frac{2K}{\epsilon} \ln \frac{|\Gamma^2| OPT^*}{t |S_t|} \leq \frac{2K}{\epsilon} \ln \frac{|\Gamma^2| OPT^*}{t} \\ & \stackrel{\textcircled{1}}{\leq} \frac{2K}{\epsilon} \ln(e + \frac{\epsilon |\Gamma^2| OPT^*}{2K}) = t, \end{aligned}$$

where $\textcircled{1}$ is due to $\frac{2K}{\epsilon} \ln(e + \frac{\epsilon OPT^* |\Gamma^2|}{2K}) \geq 2K/\epsilon$. Therefore, we can simply let $t = \frac{2K}{\epsilon} \ln(e + \frac{\epsilon OPT^* |\Gamma^2|}{2K})$ and substitute t in (7). Together with Lemma 3, (7) can be rewritten as

$$\begin{aligned} E[PR] &\geq OPT^* - 3t = OPT^* - \frac{6K}{\epsilon} \ln(e + \frac{\epsilon OPT^* |\Gamma^2|}{2K}) \\ &\geq \frac{\max_{(a,p) \in \Gamma^2} \Delta_p^a(a-p)}{K} OPT - \frac{6K}{\epsilon} \ln(e + \frac{\epsilon OPT |\Gamma^2|}{2K}) \end{aligned}$$

which ends the proof. ■

VI. EVALUATION

In this section, we provide numerical results on evaluating the performance of our DPDA scheme. Real-world dataset retrieved from Yelp [35] is adopted. The dataset is sampled from 12 metropolitan areas across 4 countries. It includes store locations, users' information, reviews, and store check-ins in the form of separate JSON files. We extract the data of city Toronto which contains 15489 businesses and more than 20000 users. In the evaluation, we treat locations of these local businesses and users as those for requestors and workers in MCS, respectively. For the rest parameters, such as task sets and workers' travel budgets, they are randomly generated. By default setting, we consider an MCS system consisting of 200 players and 80 tasks. All simulation results are the average over 100 trials.

A. Platform Revenue

Fig. 1 compares the achievable platform revenue when our DPDA scheme is implemented or not. For the latter case, we adopt the uniform pricing auction for task assignment; all $(a, p) \in \Gamma^2$ are examined to find the maximal revenue. Fig. 1(a) shows the revenue under different numbers of players when privacy budget $\epsilon = 80$ and requestor-worker-ratio $N/M = 1/4$. First of all, the revenue increases as more requestors and workers participate in the MCS market, which meets our expectation. Hence, it is desirable to design auction




Fig. 1. Platform revenue comparison with and without our DPDA scheme with respect to number of players.

schemes that can attract more participants in order to generate higher revenue for the system. Second, we observe that the revenue under our DPDA scheme is lower than that without the scheme. This is because uniform pricing auction (without differential privacy) selects the pricing pair that generates the maximal revenue, while the pricing pair in DPDA is selected following the probability (5). Apparently, the former is an upper bound of the latter. The gap between these two can be viewed as the tradeoff to protect bid privacy.

Fig. 1(b) further depicts the platform revenue when $\epsilon = 100$. Compared with Fig. 1(a), we notice that a larger privacy budget ϵ leads to a higher revenue. Specifically, the average revenue is 8.3 under $\epsilon = 80$ when there are 125 players, while it becomes 12.5 under $\epsilon = 100$. This is due to the property of pricing pair selection; with a larger ϵ , (a, p) that produces a larger $a - p$ is more likely to be selected, which leads to a higher platform revenue. Due to the same reason, the revenue is more unevenly distributed under a larger ϵ . Fig. 1(c) depicts the platform revenue when $N/M = 1$. Compared with Fig. 1(a), there are more requestors and thus more tasks, which also explains why it experiences a revenue increase when the ratio grows from $1/4$ to 1. We have a similar observation over Fig. 2, which compares platform revenue with and without DPDA with respect to number of tasks. All four figures clearly demonstrate that the revenue increases linearly in general as the task number grows.

B. Privacy Protection

To resist inference attacks, our DPDA scheme is developed based on differential privacy. Recall that its idea is to adopt probabilistic pricing rule, such that changing in a player's bid/ask does not result in significant changes of the final payment and charge profile. Thus, any player alone cannot successfully launch an inference attack through manipulating his bid/ask and analyzing the differences of the payment and




Fig. 2. Platform revenue comparison with and without our DPDA scheme with respect to number of tasks.

charge profile. Therefore, a good privacy-preserving scheme should keep changes in auction outputs as small as possible under a minor change over the auction input. To evaluate this change, we define the *privacy leakage* caused by our scheme.

Definition 7. Denote by $\wp$ and $\wp'$ pricing distributions over the set Γ^2 for differing only one element in the bid and ask profile. The privacy leakage is defined as the sum over absolute differences between the probabilities of these two pricing distributions.

$$\text{privacy leakage} = \sum_{p_l \in \wp, p'_l \in \wp'} |p'_l - p_l|.$$




Fig. 3. Impact of privacy budget ϵ to platform revenue and privacy leakage.

Fig. 3 shows the privacy leakage under different privacy budget ϵ 's. We observe that the privacy leakage increases as ϵ grows. Besides, privacy leakage keeps relatively stable when $\epsilon \leq 20$, while it increases fast after that. This property is also due to the pricing selection probability (5). Generally, with a larger ϵ , the pricing pairs that produce large difference between a and p are more likely to be selected, i.e., less diverse will be introduced into pricing; it renders our scheme more alike a pure uniform pricing auction without privacy. Meanwhile, the platform revenue also increases as ϵ grows. Its reason has been discussed in the previous section. Fig. 3 indicates that there

is a tradeoff between the platform revenue and the privacy leakage in our scheme. A larger revenue leads to a larger privacy leakage, and vice versa. Hence, the system should carefully select ϵ to strive a balance between these two.

VII. CONCLUSION

In this paper, we develop a scheme, called DPDA, to protect bid privacy of both requestors and workers in double-side markets for MCS. Our scheme is a novel combination of uniform pricing and exponential mechanism. Hence, not only can it achieve truthfulness, individual rationality and budget balance, but also guarantee differential privacy so as to resist inference attacks. Moreover, DPDA is computationally efficient and generates approximate optimal platform revenue. Both theoretical analysis and extensive simulations have been conducted to validate its performances and properties.

ACKNOWLEDGEMENT

The work of Ming Li was supported by the U.S. National Science Foundation under grants CNS-1566634 and ECCS-1711991. The work of Linke Guo was partially supported by National Science Foundation under grants ECCS-1710996 and IIS-1722731. The work of Lei Yang was supported in part by the U.S. National Science Foundation under Grants CNS-1559696 and IIA-1301726.

REFERENCES

- [1] AlgoSnap, [Online], Available: <http://crowdsignals.io/>.
- [2] Q. Zhang, Y. Wen, X. Tian, X. Gan, and X. Wang, "Incentivize crowd labeling under budget constraint," in *Proceedings of the IEEE INFOCOM*, 2015.
- [3] Z. Feng, Y. Zhu, Q. Zhang, L. M. Ni, and A. V. Vasilakos, "Trac: Truthful auction for location-aware collaborative sensing in mobile crowdsourcing," in *Proceedings of the IEEE INFOCOM*, 2014.
- [4] H. Jin, L. Su, D. Chen, K. Nahrstedt, and J. Xu, "Quality of information aware incentive mechanisms for mobile crowd sensing systems," in *Proceedings of the ACM MobiHoc*, 2015.
- [5] L. Gao, F. Hou, and J. Huang, "Providing long-term participation incentive in participatory sensing," in *Proceedings of the IEEE INFOCOM*, 2015.
- [6] Z. Duan, W. Li, and Z. Cai, "Distributed auctions for task assignment and scheduling in mobile crowdsensing systems," in *Proceedings of the IEEE ICDCS*, 2017.
- [7] C. Chen and Y. Wang, "Sparc: Strategy-proof double auction for mobile participatory sensing," in *Proceedings of the Cloud Computing and Big Data*, 2013.
- [8] X. Zhang, G. Xue, R. Yu, D. Yang, and J. Tang, "Truthful incentive mechanisms for crowdsourcing," in *Proceedings of the IEEE INFOCOM*, 2015.
- [9] H. Huang, Y. Xin, Y.-E. Sun, and W. Yang, "A truthful double auction mechanism for crowdsensing systems with max-min fairness," in *Proceedings of the IEEE WCNC*, 2017.
- [10] T. H. Hinke, H. S. Delugach, and R. P. Wolf, "Protecting databases from inference attacks," *Computers & Security*, vol. 16, no. 8, pp. 687–708, 1997.
- [11] J. Lin, D. Yang, M. Li, J. Xu, and G. Xue, "Bidguard: A framework for privacy-preserving crowdsensing incentive mechanisms," in *Proceedings of the IEEE CNS*, 2016.
- [12] H. Jin, L. Su, B. Ding, K. Nahrstedt, and N. Borisov, "Enabling privacy-preserving incentives for mobile crowd sensing systems," in *Proceedings of the IEEE ICDCS*, 2016.
- [13] Z. Huang and S. Kannan, "The exponential mechanism for social welfare: Private, truthful, and nearly optimal," in *Proceedings of the IEEE FOCS*, 2012.
- [14] T. Dimitriou and I. Krontiris, "Privacy-respecting auctions as incentive mechanisms in mobile crowd sensing," in *Proceedings of the IFIP International Conference on Information Security Theory and Practice*, 2015.
- [15] D. C. Parkes, M. O. Rabin, S. M. Shieber, and C. Thorpe, "Practical secrecy-preserving, verifiably correct and trustworthy auctions," *Electronic Commerce Research and Applications*, vol. 7, no. 3, pp. 294–312, Nov 2008.
- [16] S. Angel and M. Walfish, "Verifiable auctions for online ad exchanges," in *Proceedings of the the ACM Conference on SIGCOMM*, 2013.
- [17] M. Nojournian and D. R. Stinson, "Efficient sealed-bid auction protocols using verifiable secret sharing," in *Proceedings of the International Conference on Information Security Practice and Experience*, 2014.
- [18] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Proceedings of the IEEE FOCS*, 2007.
- [19] Y. Chen, S. Chong, I. A. Kash, T. Moran, and S. Vadhan, "Truthful mechanisms for agents that value privacy," *ACM Transactions on Economics and Computation*, vol. 4, no. 3, p. 13, Mar 2016.
- [20] D. Xiao, "Is privacy compatible with truthfulness?" in *Proceedings of the ACM Conference on Innovations in Theoretical Computer Science*, 2013.
- [21] L. Wang, D. Yang, X. Han, T. Wang, D. Zhang, and X. Ma, "Location privacy-preserving task allocation for mobile crowdsensing with differential geo-obfuscation," in *Proceedings of the International Conference on World Wide Web*, 2017.
- [22] B. Gedik and L. Liu, "Protecting location privacy with personalized k-anonymity: Architecture and algorithms," *IEEE Transactions on Mobile Computing*, vol. 7, no. 1, pp. 1–18, Jan 2008.
- [23] P. Kalnis, G. Ghinita, K. Mouratidis, and D. Papadias, "Preventing location-based identity inference in anonymous spatial queries," *IEEE Transactions on Knowledge and Data Engineering*, vol. 19, no. 12, pp. 1719–1733, Dec 2007.
- [24] J. Ni, K. Zhang, X. Lin, Q. Xia, and X. S. Shen, "Privacy-preserving mobile crowdsensing for located-based applications," in *Proceedings of the IEEE ICC*, 2017.
- [25] H. Jin, L. Su, H. Xiao, and K. Nahrstedt, "Inception: Incentivizing privacy-preserving data aggregation for mobile crowd sensing systems," in *Proceedings of the ACM MobiHoc*, 2016.
- [26] L. Blumrosen, N. Nisan, and I. Segal, "Multi-player and multi-round auctions with severely bounded communication," in *Proceedings of the European Symposium on Algorithms*, 2003.
- [27] E. Wolfstetter, "Third—and lower—price auctions," in *Beiträge zur Mikro- und zur Makroökonomik*, 2001.
- [28] A. Blum, Y. Mansour, and J. Morgenstern, "Learning valuation distributions from partial observation," in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [29] K. Deshmukh, A. V. Goldberg, J. D. Hartline, and A. R. Karlin, "Truthful and competitive double auctions," in *Proceedings of the European Symposium on Algorithms*, 2002.
- [30] C. Dwork, "Differential privacy," in *Encyclopedia of Cryptography and Security*, 2011, pp. 338–340.
- [31] A. Kothari, D. C. Parkes, and S. Suri, "Approximately-strategyproof and tractable multiunit auctions," *Decision Support Systems*, vol. 39, no. 1, pp. 105–121, Mar 2005.
- [32] A. Blum, S. Chawla, D. R. Karger, T. Lane, A. Meyerson, and M. Minkoff, "Approximation algorithms for orienteering and discounted-reward tsp," *SIAM Journal on Computing*, vol. 37, no. 2, pp. 653–670, Jun 2007.
- [33] S. He, D.-H. Shin, J. Zhang, and J. Chen, "Toward optimal allocation of location dependent tasks in crowdsensing," in *Proceedings of the IEEE INFOCOM*, 2014.
- [34] Y. S. Son, R. Baldick, K.-H. Lee, and S. Siddiqi, "Short-term electricity market auction game analysis: uniform and pay-as-bid pricing," *IEEE Transactions on Power Systems*, vol. 19, no. 4, pp. 1990–1998, 2004.
- [35] Yelp dataset challenge, [Online], Available: http://www.yelp.com/dataset_challenge.