Differentially private mechanisms for budget limited mobile crowdsourcing

Han, Kai; Liu, Huan; Tang, Shaojie; Xiao, Mingjun; Luo, Jun

2018

Han, K., Liu, H., Tang, S., Xiao, M., & Luo, J. (2019). Differentially private mechanisms for budget limited mobile crowdsourcing. IEEE Transactions on Mobile Computing, 18(4), 934-946. doi:10.1109/TMC.2018.2848265

<https://hdl.handle.net/10356/104491>

<https://doi.org/10.1109/TMC.2018.2848265>

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at:
<https://doi.org/10.1109/TMC.2018.2848265>

Downloaded on 26 Apr 2023 23:14:43 SGT

Differentially Private Mechanisms for Budget Limited Mobile Crowdsourcing

Kai Han Huan Liu Shaojie Tang Mingjun Xiao Jun Luo

Abstract—Recently, Mobile Crowdsourcing (MC) has aroused great interest on the part of both academic and industrial circles. One of the key problems in MC is designing the proper mechanisms to incentivize user participation, as users are typically self-interested and must consume a substantial amount of MC resources/costs. Although considerable research has been devoted to this problem, the majority of studies have neglected the privacy issue in mechanism design. In this study, we consider the scenario where a mobile crowdsourcing platform aims to maximize the crowdsourcing revenue under a budget constraint, and users are interested in maximizing their utility while keeping their cost private. We design differentially-private mechanisms for such a scenario under an offline setting where users bid their costs simultaneously and under an online setting where user bids are revealed one by one. We show that our mechanisms simultaneously achieve provable performance bounds with respect to several measures, including revenue, differential privacy, truthfulness and individual rationality. Finally, we also conduct extensive numerical experiments to demonstrate the effectiveness of our approach.

1 INTRODUCTION

In recent years, the spread of smart phones has led to the proliferation of Mobile Crowdsourcing (MC) applications, where collected information about interested events can be acquired by assigning MC tasks to individuals (users). Due to its wide applications, MC has already aroused great interest on the part of academic and industrial circles [1].

One of the central problems in MC is incentivizing user participation, as users are typically self-interested and incur substantial costs to perform MC tasks. However, the users may not report their true costs, as they are selfish/rational and may engage in strategic behaviour to maximize their utility. Based on these observations, many incentivization mechanisms for MC have appeared (e.g., [2]–[5]); these mechanisms usually encourage users to participate in MC and behave *truthfully* through carefully determining the monetary payments to them, while satisfying various constraints such as a predefined budget on total payments [6], [7].

Nevertheless, the majority of previous studies on incentivization mechanism design for MC have neglected another important issue—user privacy. In practice, the users’ costs for participating in MC can be important information; leaks of this information may expose users’

personal status and hence harm user utility [8]. For example, in a spatial crowdsourcing application, the users’ costs for performing MC tasks could be determined by their traveling costs to the Point of Interests (PoIs) [9], so having the cost information in hand would allow attackers to infer the location information about the users [10].

As the current incentivization mechanisms for crowdsourcing determine the payments to the users based on the users’ costs¹, the cost information of the users can be easily leaked to any third party (or adversary) who observes the payment profile calculated by the crowdsourcing mechanism. One possible way to address this problem is to use the traditional syntactic approaches such as k -anonymity and ℓ -diversity [12], [13] to protect the users’ privacy. Roughly speaking, these syntactic approaches try to generalize the data entries such that the ability of an adversary to link a “quasi-identifier” tuple to sensitive values is restricted. However, it has been proved that these approaches are vulnerable, especially when the adversary has strong background knowledge [14]. Therefore, a more prevailing approach for data privacy adopted by the recent work is Differential Privacy (DP) [15]. The basic idea of DP is to add noises to the answers of data queries such that it becomes harder for an adversary who observes the output of the algorithm to distinguish two neighbouring input datasets of the algorithm (in a probabilistic sense) [15]. Compared to the traditional syntactic approaches, DP has a more rigorous mathematical framework for defining and preserving privacy, and it also adopts a stronger model on adversary’s background knowledge, i.e., the

1. This is based on the rationale of “individual rationality” or “voluntary participation” [11], i.e., no user who truthfully participates in the crowdsourcing campaign should be paid less than her/his private cost.

- Kai Han and Huan Liu are with School of Computer Science and Technology / Suzhou Institute for Advanced Study, University of Science and Technology of China, China. E-mail: hankai@ustc.edu.cn
- Shaojie Tang is with Naveen Jindal School of Management, University of Texas at Dallas, United States. E-mail: tangshaojie@gmail.com
- Mingjun Xiao is with School of Computer Science and Technology, University of Science and Technology of China, China. E-mail: xiaomj@ustc.edu.cn
- Jun Luo is with School of Computer Science and Engineering, Nanyang Technological University, Singapore. E-mail: junluo@ntu.edu.sg

adversary can observe all the data records except for the one whose privacy is to be protected.

However, designing differentially private incentivization mechanisms for MC is complex, as we typically need to simultaneously achieve other performance goals except DP (e.g., revenue maximization) under various system constraints (e.g., the budget constraint for payments). Indeed, the desired goals even appear to be self-contradictory. For example, it can be easily understood that payments to users must be sensitive to the adjustment of any single user's cost to optimize the system's revenue under a limited budget, but the concern about DP requires the mechanism to be insensitive to such an adjustment (to protect the user privacy). In addition, the truthfulness problem must be addressed at the same time; otherwise, the user costs revealed to the system can be noisy (and strategic) data even if DP is not taken into account. In summary, these entangled issues make it extremely difficult to design privacy-preserving incentivization mechanisms for MC, and a good mechanism should strive to seek a balance between truthfulness, privacy and other optimization goals such as revenue maximization under rigorous resource constraints.

Due to the difficulties described above, the existing incentivization mechanisms for mobile crowdsourcing usually drop one or more requirements among truthfulness, budget-feasibility and differential privacy, so it is simpler for them to get some provable performance bounds. For example, a large body of existing studies on crowdsourcing incentivization (e.g., [6], [7], [16], [17]) have neglected the privacy issue, while some other work considering DP has neglected the budget constraint [8], [18]. To the best of our knowledge, no previous work has proposed a truthful and differentially private incentivization mechanism for mobile crowdsourcing under a budget constraint.

1.1 Our Contributions

In this paper, we study the problem of maximizing crowdsourcing revenue under a budget constraint on payments to users and propose mechanisms that achieve budget-feasibility, truthfulness, differential privacy and high revenue simultaneously. We consider both the offline setting where all users show up simultaneously and the online setting where users arrive sequentially in an arbitrary order. More specifically, our contributions can be summarized as follows:

- 1) In the offline setting, we first propose a benchmark mechanism called PWDP with a $\frac{1}{2}$ performance ratio on the revenue without considering DP, and then provide a mechanism achieving ϵ -differential privacy and a performance ratio close to that of PWDP.
- 2) In the online setting, we propose DPP-UCB, a dynamic pricing mechanism based on the multi-armed-bandit paradigm [19]. We prove that DPP-UCB achieves ϵ -differential privacy and an

$\mathcal{O}(\log W \log \log W)$ regret bound with respect to revenue, where W is the predefined budget limit.

- 3) In addition to DP, we prove that all our proposed mechanisms achieve other nice properties including truthfulness, budget-feasibility and individual rationality [11].
- 4) We conduct extensive numerical experiments to compare our algorithms with related studies, and the experimental results demonstrate the effectiveness of our approach.
- 5) To the best of our knowledge, we are the first to propose differentially private and budget-limited mechanisms for mobile crowdsourcing with provable performance bounds, both under the offline setting and under the online setting.

The rest of our paper is organized as follows. We first formally formulate our problem in Sec. 2, and then introduce our mechanisms as well as their performance analysis in Sec. 3 and Sec. 4. The experimental results are shown in Sec. 5. We discuss related work in Sec. 6 before concluding the paper in Sec. 7.

2 PROBLEM FORMULATION

We assume that there are a mobile crowdsourcing platform and a set of users in $[m] \triangleq \{1, 2, \dots, m\}$. Each user $j \in [m]$ has a private cost $c_j \in [\theta_1, \theta_2]$ for performing one crowdsourcing task, where θ_1, θ_2 are known constants. Following [6], [8], [17], [20], we assume that the platform has a budget W for paying the users, and the monetary payment to any user is selected from a set $\mathcal{S}$ of candidate prices. Note that the users' costs and the rewards to them are usually monetized in real crowdsourcing applications. For example, the crowdsourcing tasks in Amazon's Mechanical Turk [21] are usually priced at several cents/dollars. Based on this observation, we assume that $\mathcal{S}$ is a discrete set, which is also assumed in some related work such as [22]. Therefore, the set $\mathcal{S}$ can be represented by $\{s_1, s_2, \dots, s_k\}$, where $s_1 < s_2 < \dots < s_k$.

A crowdsourcing mechanism (possibly a randomized algorithm) of the platform finds a payment profile $\mathbf{p} = \langle p_1, p_2, \dots, p_m \rangle \in \mathcal{S}^m$ and a winner set $\mathcal{N} \subseteq [m]$ under the constraint of $\sum_{j \in \mathcal{N}} p_j \leq W$, such that any user $j \in [m]$ is assigned one crowdsourcing task (and receives payment p_j) if and only if $j \in \mathcal{N}$. The utility of any user $j \in [m]$ in the crowdsourcing mechanism can be written as

$$u_j = \begin{cases} p_j - c_j; & \text{if } j \in \mathcal{N} \\ 0; & \text{otherwise} \end{cases}$$

We also assume that completing the crowdsourcing tasks yields identical values, which has also been assumed by some related work such as [17], [20], [22]. In practice, this assumption holds in crowdsourcing applications with homogeneous tasks (e.g., reCAPTCHA [23]). Based on this assumption, the revenue of our crowdsourcing mechanism is defined as the total (expected) number of crowdsourcing tasks assigned to the users, i.e., $\mathbb{E}\{|\mathcal{N}|\}$.

TABLE 1: Some Frequently Used Notations

Notation	Description
m	The number of crowdsourcing users
$[m]$	The set $\{1, \dots, m\}$
W	The budget for paying the users
$\mathcal{N}$	The set of users who are assigned tasks
$\mathcal{S}$	The set of candidate prices $\{s_1, \dots, s_k\}$
k	The cardinality of $\mathcal{S}$
$\mathbf{x}$	Vector $(x_1, \dots, x_m) \in \{0, 1\}^m; \forall j \in [m] : x_j = 1 \Leftrightarrow j \in \mathcal{N}$
$\mathbf{p}$	Vector $(p_1, \dots, p_m); p_j$ is the payment to user j
$\mathbf{c}$	Vector $(c_1, \dots, c_m); c_j$ is the cost of user j
$\mathbf{b}$	Vector $(b_1, \dots, b_m); b_j$ is the bid of user j under the offline setting
$\xi(b_j)$	$\min\{e e \in \mathcal{S} \wedge b_j \leq e\}$
$r(\mathbf{b}, e)$	$\min\{W/e, f(\mathbf{b}, e)\}$
$f(\mathbf{b}, e)$	$ \{j j \in [m] \wedge \xi(b_j) \leq e\} $
D_l	The probability that any user's cost is no more than s_l under the online setting
φ_l	$\min\{mD_l, W/s_l\}$
$n_{l,t}$	The total number of times that price s_l has been posted until time t under the online setting
$\mathcal{H}_{l,t}$	$\frac{\sqrt{8 \log 4t^4}}{\epsilon n_{l,t}} (1 + \log n_{l,t})$
$\sigma_{l,t}$	$\sqrt{5 \ln t / 2n_{l,t}}$
$\mathcal{D}_{l,t}$	The estimation on D_l at any time t by using the hybrid mechanism under the online setting
$\tilde{\varphi}_{l,t}$	$\min\{m(\mathcal{D}_{l,t} + \sigma_{l,t} + \mathcal{H}_{l,t}), W/s_l\}$
$[\theta_1, \theta_2]$	The range of users' costs

If we know all users' true costs, then the problem of maximizing the revenue under the budget W can be optimally solved in polynomial time. Specifically, we can first sort the users according to the non-decreasing order of their costs, and then select the users according to this order until the total cost exceeds W . However, the costs of the users are usually unknown, which makes our problem much more complex. To address this problem, a possible approach is to solicit the cost information from the users such that the crowdsourcing revenue can be optimized. However, as the users are selfish and rational, they may report false information to the platform to maximize their own utilities. Therefore, a crowdsourcing mechanism should encourage the users to tell the truth, by aligning the personal interests of the users with the system goal of revenue maximization. Moreover, it should also guarantee that no user gets a negative utility as long as the user behaves honestly. Based on these considerations, we formally introduce the definitions on *truthfulness* and *Individual Rationality* (IR):

Definition 1. A crowdsourcing mechanism is called *dominant strategy truthful* iff any user $j \in [m]$ maximizes her/his utility u_j by reporting truthfully to the platform, regardless of how the other users report. The mechanism is called *individually rational* iff any truthful user gets a non-negative utility.

Besides achieving truthfulness, we also aim to protect differential privacy of the users with respect to their costs for performing the crowdsourcing tasks. Generally, protecting the privacy against an attacker with more background knowledge would be harder than that against

Fig. 1: User bidding under the offline model

an attacker with less background knowledge [14], [15]. In this paper, we assume that the attacker has strong background knowledge, i.e., she/he can observe the payment profile $\mathbf{p}$. Note that the attacker can either watch $\mathbf{p}$ when it is published by the mechanism, or observe it through some other ways (e.g., somehow snooping on the paying records of the mechanism) when $\mathbf{p}$ is not publicly published. It would also be very interesting to consider the case where the attacker has no background knowledge (i.e., she/he cannot get the information of the payment profile), but this problem might be out of the scope of our paper due to the space constraint, and we plan to study it in our future work.

We will also consider two models that address how users show up in the system, i.e., the offline model and the online model. Although we have the same goal of maximizing crowdsourcing revenue under both of these models, the definitions on other measures such as differential privacy are different. In the sequel, we explain these models in detail, and we also list some frequently used notations in Table 1.

2.1 Offline Model

In the offline model, we assume that all users appear at the same time with arbitrary costs in $[\theta_1, \theta_2]$. The crowdsourcing mechanism reveals users' private costs by soliciting a bidding vector $\mathbf{b} = \langle b_1, \dots, b_m \rangle$ from users, where b_j is the reported cost of user j for any $j \in [m]$, and then decides the payment profile $\mathbf{p}$ as well as the winner set $\mathcal{N}$. An intuitive illustration for the offline model is shown in Fig. 1.

As discussed in Sec. 1, the payment profile $\mathbf{p}$ should be "insensitive" to changes of a single user's cost/bid, so the privacy of users can be protected. Based on this idea and [15], we provide the following definition of privacy:

Definition 2. (ϵ -differential privacy for the offline model) A crowdsourcing mechanism for the offline model is called ϵ -differentially private iff for any $j \in [m]$, any $\mathcal{P} \subseteq \mathcal{S}^m$ and any two bidding vectors $\mathbf{b} = \langle b_j, b_{-j} \rangle$ and $\mathbf{b}' = \langle b'_j, b_{-j} \rangle$, we have $\mathbb{P}\{\mathbf{p} \in \mathcal{P} | \langle b_j, b_{-j} \rangle\} \leq \exp(\epsilon) \mathbb{P}\{\mathbf{p} \in \mathcal{P} | \langle b'_j, b_{-j} \rangle\}$, where b_{-j} denotes the bids of the users in $[m] \setminus \{j\}$.

Example 1: Suppose that a pricing mechanism selects a payment uniformly at random from $\mathcal{S}$ for each user. Although such a pricing mechanism does not guarantee individual rationality, it satisfies ϵ -differential privacy according to Definition 2. Indeed, we have $\epsilon = 0$ in this

case, as the variations of the users' bids do not affect the payments at all.

It is well acknowledged that randomized mechanisms should be designed to guarantee differential privacy [15]. Therefore, the concept of truthfulness in Definition 1 can be relaxed to the following definition on truthfulness-in-expectation:

Definition 3. (*truthfulness in expectation*) A crowdsourcing mechanism is called ϵ -approximate truthful-in-expectation iff $\mathbb{E}\{u_j | \langle c_j, b_{-j} \rangle\} \geq \mathbb{E}\{u_j | \langle b_j, b_{-j} \rangle\} - \epsilon$ holds for any j, b_j and b_{-j} . When $\epsilon \leq 0$, the mechanism is called exactly truthful-in-expectation.

2.2 Online Model

In the online case, we assume that users arrive one by one in an arbitrary order, and their private costs are drawn independently from an unknown distribution. This assumption generalizes the standard Bayesian mechanism design model adopted in [4], [5], [11], where the cost distribution of users is assumed to be known in advance.

When any user $j \in [m]$ shows up, the mechanism selects and posts a take-it-or-leave-it price $p_j \in \mathcal{S}$ for user j based on the prices selected in the history and past observations on users $\{1, \dots, j-1\}$. After observing the posted price p_j , user j reports $x_j \in \{0, 1\}$ to the mechanism, where x_j denotes whether user j accepts the price ($x_j = 1$ for acceptance). The mechanism adds j to the winner set $\mathcal{N}$ if and only if $x_j = 1$. This process continues until either the budget W is depleted or the last user m has been processed. An intuitive illustration for posted pricing under the online model is shown in Fig. 2.

Note that the posted-pricing scheme described above has also been adopted in many commercial crowdsourcing systems such as Amazon's Mechanical Turk. However, achieving differential privacy is trickier in the online model than that in the offline model, as we have to guarantee that the mechanism achieves differential privacy at each time a user shows up. Therefore, we revise Definition 2 and get the following definition on DP for the online model:

Definition 4. (ϵ -differential privacy for the online model) Let $\mathbf{p}_j$ and $\mathbf{x}_j$ denote $\langle p_1, \dots, p_j \rangle$ and $\langle x_1, \dots, x_j \rangle$, respectively. Any pair $(\mathbf{x}_j, \mathbf{x}'_j) \in \{0, 1\}^j \times \{0, 1\}^j$ is called adjacent iff $\mathbf{x}_j$ and $\mathbf{x}'_j$ differ in at most one element. An online pricing mechanism is ϵ -differentially private iff for any $j \in [m]$, any $P \subseteq \mathcal{S}$, any $\mathbf{p}_{j-1} \in \mathcal{S}^{j-1}$ and any adjacent pair $(\mathbf{x}_{j-1}, \mathbf{x}'_{j-1}) \in \{0, 1\}^{j-1} \times \{0, 1\}^{j-1}$, we have $\mathbb{P}\{p_j \in P | \mathbf{p}_{j-1}, \mathbf{x}_{j-1}\} \leq \exp(\epsilon) \mathbb{P}\{p_j \in P | \mathbf{p}_{j-1}, \mathbf{x}'_{j-1}\}$.

Example 2: Suppose that $j = 3$. Then the pair $(\mathbf{x}_{j-1}, \mathbf{x}'_{j-1}) = (\langle 0, 1 \rangle, \langle 0, 0 \rangle)$ is adjacent. However, if $(\mathbf{x}_{j-1}, \mathbf{x}'_{j-1}) = (\langle 0, 0 \rangle, \langle 1, 1 \rangle)$, then $(\mathbf{x}_{j-1}, \mathbf{x}'_{j-1})$ is not adjacent. Indeed, there are totally 12 possible adjacent pairs for $(\mathbf{x}_{j-1}, \mathbf{x}'_{j-1})$. If none of these adjacent pairs

Fig. 2: Posted pricing under the online model

result in a significant change of p_j (in the probabilistic sense explained in Definition 4), then such a pricing mechanism achieves differential privacy (at time j). Indeed, the conditions in Definition 4 are stronger, as it requires that the pricing mechanism achieves differential privacy for any $1 \leq j \leq m$.

3 MECHANISM DESIGN FOR THE OFFLINE MODEL

In this section, we design mechanisms for the offline model. We first introduce a benchmark algorithm without considering DP. Then, we propose our privacy-preserving mechanism and analyze its performance.

3.1 A Truthful Mechanism without Considering DP

It can be seen that the revenue of any mechanism that achieves individual rationality under the offline model can be upper-bounded by

$$R_{opt} = \max \left\{ |U| : U \in 2^{[m]} \wedge \sum_{j \in U} \xi(b_j) \leq W \right\}, \quad (1)$$

where $\xi(b_j) = \min\{e | e \in \mathcal{S} \wedge b_j \leq e\}$. Moreover, we can optimally find a set of users $U \subseteq [m]$ such that the revenue got from U is R_{opt} , assuming that all users in $[m]$ bid truthfully. However, when users act strategically, we have to sacrifice some revenue to ensure truthfulness. A naive idea for designing a truthful mechanism is to use the VCG mechanism [11], but it is known that this approach can compromise budget-feasibility [2]. Therefore, we propose PWDP, an algorithm that simultaneously achieves truthfulness and budget-feasibility (but without achieving DP), as shown by Algorithm 1. Algorithm 1 serves as a benchmark algorithm for our problem under the offline setting.

The idea of Algorithm 1 is explained as follows. We first sort the users according to the non-decreasing order of the minimum payments to them to ensure individual rationality (line 1), and then add the users into the winner set according to their orders sorted by line 1 and a filter condition in line 3. After that, we calculate the payment profile to the users (lines 5-7), such that each user $j \in \mathcal{N}$ is equally paid with a "threshold value" to guarantee truthfulness. More specifically, we use K to denote the maximum equal payment to each user in $\mathcal{N}$ under the budget constraint (line 5); and the payment p_j to any winner $j \in \mathcal{N}$ is set to the smaller one between K and $\xi(b_{d_{q+1}})$ (line 7), where q is the cardinality of $\mathcal{N}$. Intuitively, such a payment profile guarantees that, if any of the winner $j \in \mathcal{N}$ unilaterally raises her/his bid to be larger than p_j , then she/he

Algorithm 1: The PWDP algorithm

Input: $\mathbf{b}, W, m, \mathcal{S}$

Output: $\mathbf{p}, \mathcal{N}$

```

1 Sort the users in  $[m]$  into  $d_1, \dots, d_m$  such that
    $\xi(b_{d_1}) \leq \xi(b_{d_2}) \leq \dots \leq \xi(b_{d_m})$ , ties broken according
   to the id of the users;
2 for  $j = m$  to 1 do
3   if  $\xi(b_{d_j}) \leq W/j$  then
4      $\mathcal{N} \leftarrow \{d_1, \dots, d_j\}$ ; break;
5  $q \leftarrow |\mathcal{N}|$ ;  $\mathbf{p} \leftarrow \mathbf{0}$ ;  $K \leftarrow \max\{e | e \in \mathcal{S} \wedge e \leq W/q\}$ 
6 for  $j = 1$  to  $m$  do
7   if  $j \in \mathcal{N}$  then  $p_j \leftarrow \min\{\xi(b_{d_{q+1}}), K\}$ ;
8 return  $\mathbf{p}, \mathcal{N}$ 

```

would no longer be selected as a winner. This property is sufficient for proving the truthfulness of Algorithm 1 due to the Myerson's Lemma [11], as shown by the following theorem:

Theorem 1. *Algorithm 1 achieves dominant strategy truthfulness, individual rationality and budget feasibility.*

Proof: Clearly, Algorithm 1 achieves budget feasibility according to line 7. Let $\mathcal{N}(\mathbf{b})$ denote the winner set output by Algorithm 1 given the input bidding vector $\mathbf{b}$. To prove the truthfulness, we only need to prove the following conditions according to the celebrated Myerson's lemma [11]: (1) Algorithm 1 is monotone, i.e., $\forall j \in [m]$, if $b'_j \leq b_j$; then, $j \in \mathcal{N}(b_j, b_{-j})$ implies $j \in \mathcal{N}(b'_j, b_{-j})$ for every b_{-j} ; (2) Each winner is paid the threshold value, i.e., each winner j is paid $p_j = \inf\{b_j : j \notin \mathcal{N}(b_j, b_{-j})\}$. Clearly, condition (1) holds. Next, we prove condition (2):

Case 1: $\xi(b_{d_{q+1}}) \leq K$:

In this case, any user $j \in \mathcal{N}$ has utility $u_j = \xi(b_{d_{q+1}}) - c_j \geq \xi(b_j) - c_j$, which is non-negative if the user is truthful (i.e., $b_j = c_j$). In addition, any user $j \in \mathcal{N}$ bidding any $b'_j \leq \xi(b_{d_{q+1}})$ still wins, as she/he is still ranked before user $b_{d_{q+1}}$ with such a bid b'_j , and we have $\xi(b'_j) \leq \xi(b_{d_{q+1}}) \leq W/q \leq W/z$ for any $z \in [q]$. However, the user will lose by bidding any $b'_j > \xi(b_{d_{q+1}})$, because she/he would be ranked after $b_{d_{q+1}}$ with such a bid, and we have $\xi(b'_j) \geq \xi(b_{d_{q+1}}) > W/(q+1) \geq W/z$ for any $z \geq q+1$.

Case 2: $K < \xi(b_{d_{q+1}})$:

In this case, any user $j \in \mathcal{N}$ has the utility $u_j = K - c_j \geq \xi(b_{d_q}) - c_j \geq \xi(b_j) - c_j$, which is non-negative if $b_j = c_j$. Any user $j \in \mathcal{N}$ bidding any $b'_j \leq K$ still wins, as $\xi(b'_j) < \xi(b_{d_{q+1}})$ and $\xi(b'_j) \leq W/q \leq W/z$ for any $z \in [q]$. Bidding $b'_j > K$ will make user j lose, as $\xi(b'_j) > W/q \geq \xi(b_{d_q}) \geq W/z$ for any $z \geq q$.

Note that the above reasoning also proves individual rationality. Hence, the theorem follows. $\square$

Besides guaranteeing truthfulness, Algorithm 1 also achieves a constant performance ratio on the revenue:

Theorem 2. *The revenue obtained by Algorithm 1 is at least*

$$\frac{1}{2} R_{opt}.$$

Proof: Suppose that Algorithm 1 outputs $\mathcal{N} = \{d_1, \dots, d_q\}$ and $q_{max} = \max\{z | \sum_{t=1}^z \xi(b_{d_t}) \leq W\}$. So $R_{opt} = q_{max}$ and $q \leq q_{max}$. If $q = q_{max}$, the theorem trivially holds. Thus, we assume $q < q_{max}$. Note that we have $\xi(b_{d_{q+1}}) > W/(q+1)$ and $(q_{max} - q)\xi(b_{d_{q+1}}) \leq \sum_{t=1}^{q_{max}} \xi(b_{d_t}) \leq W$, so we get $q_{max} \leq 2q$. Hence, the theorem follows. $\square$

Example 3: Suppose that $\mathcal{S} = \{1, \dots, 10\}$, $W = 11$, $m = 5$ and $\mathbf{b} = \langle b_1, b_2, b_3, b_4, b_5 \rangle = \langle 2, 5, 1, 3, 6 \rangle$. According to line 1 of Algorithm 1, we have $\langle d_1, d_2, d_3, d_4, d_5 \rangle = \langle 3, 1, 4, 2, 5 \rangle$. According to lines 2-4 of Algorithm 1, we have $\mathcal{N} = \{1, 3, 4\}$. According to lines 5-7 of Algorithm 1, we have $\langle p_1, p_2, p_3, p_4, p_5 \rangle = \langle 3, 0, 3, 3, 0 \rangle$. It is noted that, when all the users are truthful, an optimal solution is to select the users 1, 2, 3 and 4, and pay them 2, 5, 1, 3, respectively. This optimal solution has the revenue 4, while the solution output by Algorithm 1 has the revenue 3. Therefore, Algorithm 1 achieves at least half of the optimal revenue in this instance of our problem.

3.2 The OPEX Mechanism

In this section, we propose a privacy-preserving mechanism called OPEX, shown in Algorithm 2. The idea of Algorithm 2 is that we first randomly select a payment profile with identical payments to all users (line 1) and then determine the winner set using this payment profile and the budget constraint W (lines 2-4). To optimize revenue, OPEX selects each payment profile $\langle p_{ox}, p_{ox}, \dots, p_{ox} \rangle$ with the probability proportional to $r(\mathbf{b}, p_{ox})$, which denotes the revenue that can be gained by using this payment profile. More specifically, the function r is defined as

$$\forall e \in \mathcal{S} : r(\mathbf{b}, e) = \min\{\lfloor W/e \rfloor, f(\mathbf{b}, e)\}, \quad (2)$$

where $f(\mathbf{b}, e) \triangleq |\{j | j \in [m] \wedge \xi(b_j) \leq e\}|$.

According to lines 2-4 of Algorithm 2, any bidder j selected in the winner set L satisfies the condition $b_j \leq p_{ox}$. Therefore, the utility of any truthful bidder j is either 0 or $p_{ox} - b_j \geq 0$. Thus, OPEX achieves individual rationality. In addition, from lines 3-4, we can see that the total payment of OPEX is no more than $\ell p_{ox} \leq \lfloor W/p_{ox} \rfloor p_{ox} \leq W$, so OPEX achieves budget-feasibility. Moreover, inspired by the performance analysis for the exponential mechanism [15], [24], we can also prove:

Theorem 3. *OPEX achieves ϵ -differential privacy and 2ϵ -approximate truthfulness in expectation.*

Proof: For any $\mathbf{b} = \langle b_i, b_{-i} \rangle$, $\mathbf{b}' = \langle b'_i, b_{-i} \rangle$ and any $s \in \mathcal{S}$, we have

$$\begin{aligned}
& \mathbb{P}\{\mathbf{p} = \langle s, \dots, s \rangle | \mathbf{b}\} / \mathbb{P}\{\mathbf{p} = \langle s, \dots, s \rangle | \mathbf{b}'\} \\
&= \frac{\exp\{\epsilon r(\mathbf{b}, s)/2\}}{\sum_{e \in \mathcal{S}} \exp\{\epsilon r(\mathbf{b}, e)/2\}} \cdot \frac{\sum_{e \in \mathcal{S}} \exp\{\epsilon r(\mathbf{b}', e)/2\}}{\exp\{\epsilon r(\mathbf{b}', s)/2\}} \\
&\leq \exp\{\epsilon \Delta_r/2\} \exp\{\epsilon \Delta_r/2\} \\
&\leq \exp\{\epsilon \Delta_r\}, \quad (3)
\end{aligned}$$

Algorithm 2: The OPEX mechanism

Input: $\mathbf{b}, W, m, \mathcal{S}, \epsilon \in [0, 1]$

Output: $\mathbf{p}, \mathcal{N}$

- 1 Select $p_{ox} \in \mathcal{S}$ with probability proportional to $\exp\{\epsilon r(\mathbf{b}, p_{ox})/2\}$;
 - 2 $K \leftarrow \{j | j \in [m] \wedge b_j \leq p_{ox}\}$;
 - 3 $\ell \leftarrow \min\{|K|, \lfloor W/p_{ox} \rfloor\}$;
 - 4 $\mathcal{N} \leftarrow$ an arbitrary set $L \subseteq K$ such that $|L| = \ell$;
 - 5 **return** $\mathcal{N}, \mathbf{p} = \langle p_{ox}, p_{ox}, \dots, p_{ox} \rangle$
-

where $\Delta r = \max_{e \in \mathcal{S}} |r(\mathbf{b}, e) - r(\mathbf{b}', e)|$.

Note that $\forall e \in \mathcal{S} : |f(\mathbf{b}, e) - f(\mathbf{b}', e)| \leq 1$. When $\lfloor W/e \rfloor \leq \min\{f(\mathbf{b}, e), f(\mathbf{b}', e)\}$, we have $|r(\mathbf{b}, e) - r(\mathbf{b}', e)| = 0$. When $\lfloor W/e \rfloor \geq \max\{f(\mathbf{b}, e), f(\mathbf{b}', e)\}$, we have $|r(\mathbf{b}, e) - r(\mathbf{b}', e)| \leq 1$. When $f(\mathbf{b}, e) \leq \lfloor W/e \rfloor \leq f(\mathbf{b}', e)$ or $f(\mathbf{b}', e) \leq \lfloor W/e \rfloor \leq f(\mathbf{b}, e)$, we have $|r(\mathbf{b}, e) - r(\mathbf{b}', e)| \leq |f(\mathbf{b}, e) - f(\mathbf{b}', e)| \leq 1$. Therefore, we have $\Delta r \leq 1$, and hence OPEX achieves ϵ -differential privacy. Finally, according to Proposition 10.1 of [15], it can be seen that OPEX achieves 2ϵ approximate truthfulness truthful-in-expectation. $\square$

Example 4: Reconsider the problem instance in Example 3. We have $f(\mathbf{b}, 1) = 1, f(\mathbf{b}, 2) = 2, f(\mathbf{b}, 3) = f(\mathbf{b}, 4) = 3, f(\mathbf{b}, 5) = 4$ and $f(\mathbf{b}, i) = 5$ for $i \geq 6$. According to Eqn. (2), we have $r(\mathbf{b}, 1) = 1, r(\mathbf{b}, 2) = 2, r(\mathbf{b}, 3) = 3, r(\mathbf{b}, 4) = r(\mathbf{b}, 5) = 2$ and $r(\mathbf{b}, i) = 1$ for $i \geq 6$. Therefore, according to line 1 of Algorithm 2, the probability that Algorithm 2 selects $p_{ox} = 3$ is the largest among all the ten prices in $\mathcal{S}$. When $p_{ox} = 3$, we have $K = \{1, 3, 4\}$ and $\ell = 3$ according to lines 3-4 of Algorithm 2, so we have $\mathcal{N} = \{1, 3, 4\}$ according to line 4 of Algorithm 2.

Next, we analyze the performance of OPEX in terms of revenue. Note that equation (1) implies that users should be paid unequally to obtain the maximal revenue R_{opt} . However, the following theorem reveals that we can obtain at least half of the optimal revenue by using identical payments to all users:

Theorem 4. Let $s_{opt} = \arg \max_{e \in \mathcal{S}} r(\mathbf{b}, e)$. Then, we have $R_{opt} \leq 2r(\mathbf{b}, s_{opt})$.

Proof: We assume w.l.o.g. that $\xi(b_1) \leq \xi(b_2) \leq \dots \leq \xi(b_m)$. Let $v = r(\mathbf{b}, s_{opt})$ and $h = R_{opt}$. If $v = m$ or $v \geq h$, the theorem is trivially proved. When $v < m$ and $v < h$, we must have $v\xi(b_v) \leq W$ and $(v+1)\xi(b_{v+1}) > W$. Note that we also have $\sum_{j=1}^h \xi(b_j) \leq W$. Therefore, we get $W \geq \sum_{j=v+1}^h \xi(b_j) \geq (h-v)\xi(b_{v+1}) > \frac{h-v}{v+1}W$, so $R_{opt} = h \leq 2v = 2r(\mathbf{b}, s_{opt})$. $\square$

Note that OPEX outputs a payment profile with identical payments. So, we can use Theorem 4 to prove Theorem 5:

Theorem 5. We have $r(\mathbf{b}, p_{ox}) \geq \frac{1}{2}R_{opt} - \mathcal{O}(\frac{\log k}{\epsilon})$ with high probability.

Proof: Let $\vartheta = \frac{1}{2}R_{opt} - \frac{2 \ln k + z}{\epsilon}$, where z is a positive constant (e.g., we can set $z = 10$). Using Theorem 4, we

have:

$$\begin{aligned} \mathbb{P}\{r(\mathbf{b}, p_{ox}) < \vartheta\} &\leq \frac{k \exp\{\epsilon \vartheta / 2\}}{\exp\{\epsilon r(\mathbf{b}, s_{opt}) / 2\}} \\ &\leq k \exp\{\epsilon \vartheta / 2\} / \exp\{\epsilon R_{opt} / 4\} \\ &= k \exp\{\epsilon(2\vartheta - R_{opt}) / 4\} = \exp\{-z/2\}, \end{aligned} \quad (4)$$

which implies that $\mathbb{P}\{r(\mathbf{b}, p_{ox}) \geq \vartheta\} > 1 - \exp\{-z/2\}$. Hence, the theorem follows. $\square$

It can be seen that the only difference between the bounds shown in Theorem 2 and Theorem 5 is the $\mathcal{O}(\frac{\log k}{\epsilon})$ additive factor. This suggests that we would not lose too much revenue by considering differential privacy.

4 MECHANISM DESIGN FOR THE ONLINE MODEL

In the online case, our problem becomes more difficult, as we have to guarantee DP at any time, and the cost distribution of the users is unknown. Based on these considerations, the intuitive idea of our algorithm can be roughly explained as follows. Instead of using the true means of the users' costs (which are unknown), we try to use the empirical means to design a pricing mechanism. As such, if the empirical means are accurate enough, then the revenue got by our pricing mechanism should be close to the optimal revenue got under the Bayesian setting. However, directly using the empirical means could breach the users' privacy. Therefore, we will add noises into the empirical means before we use them, and carefully adjust the "noise level" such that the resulted noisy empirical means can achieve high accuracy while satisfying differential privacy.

Based on the above idea, we will leverage the multi-armed-bandit (MAB) paradigm [25] and the hybrid mechanism [26] to design our algorithm. Intuitively, the MAB paradigm enables us to quickly find an approximate solution based on the empirical means, while the hybrid mechanism [26] provides a method for calculating noisy empirical means that achieve differential privacy and high accuracy. However, it is highly non-trivial to design an algorithm with provable performance bounds based on these intuitive ideas, and we will elaborate our algorithms in the following sections.

4.1 The Hybrid Mechanism

The goal of the hybrid mechanism [26] is to publish the aggregation of private data. More specifically, suppose that there is a data sequence $d_1, d_2, \dots$ where $d_n \in \{0, 1\}$ for any $n \in \mathbb{N}$. The hybrid mechanism is a function B satisfying $B(\{d_1, \dots, d_n\}) = \sum_{j \in [n]} d_j + \gamma_n$ for any $n \in \mathbb{N}$, where γ_n is a random variable (i.e., the "noise"). By carefully choosing the probability distribution of γ_n , the hybrid mechanism can achieve ϵ -differential privacy, i.e., for any $n \in \mathbb{N}$, any $j \in [n]$, any $A = \{d_1, \dots, d_j, \dots, d_n\}$, $A' = \{d_1, \dots, d'_j, \dots, d_n\}$ and any $Z \subseteq \mathbb{R}$, we have $\mathbb{P}\{B(A) \in Z\} \leq \exp(\epsilon) \mathbb{P}\{B(A') \in Z\}$. Moreover, it

has been shown that the hybrid mechanism has the following nice property:

Lemma 1 ([26], [27]). Define $\chi(\kappa, n) = \frac{\sqrt{8}}{\epsilon} \log\left(\frac{4}{\kappa}\right) \log n + \frac{\sqrt{8}}{\epsilon} \log\left(\frac{4}{\kappa}\right)$. For any $\kappa \leq n^{-d}$ ($d > 0$) and any $n \in \mathbb{N}$, we have $\mathbb{P}\{|\gamma_n| \geq \chi(\kappa, n)\} \leq \kappa$.

Intuitively, Lemma 1 implies that we do not have to add a too large “noise” (i.e., γ_n) to the sum of $d_1, \dots, d_n$ such that the noisy sum $B(\{d_1, \dots, d_n\})$ satisfies ϵ -differential privacy. More specifically, as indicated by Lemma 1, γ_n only needs to be logarithmic to n in a probabilistic sense. This property will be useful for the performance analysis of our algorithms in Sec. 4.3.

4.2 The DPP-UCB Algorithm

With the hybrid mechanism, we propose a privacy preserving mechanism called DPP-UCB (shown in Algorithm 3). Clearly, if we consider a Bayesian setting where the cost distribution of the users is known, then the revenue got by offering any price $s_l \in \mathcal{S}$ to all users is at most $\varphi_l = \min\{mD_l, \frac{W}{s_l}\}$, where $D_l = \mathbb{P}\{c_1 \leq s_l\}$. Moreover, it has been known that, the maximum revenue that we can get by offering a single price to all users is a constant approximation to the revenue got by any truthful mechanism [2], [28]. Therefore, the main idea of DPP-UCB is trying to select the price s_{l^*} at each time t , where $l^* = \arg \max_{l \in [|\mathcal{S}|]} \varphi_l$.

However, D_l and φ_l are unknown for any $l \in [|\mathcal{S}|]$. Therefore, we use the hybrid mechanism to get a noisy empirical mean $\mathcal{D}_{l,t}$ as the estimation of D_l for any $l \in [|\mathcal{S}|]$ and $t > 0$. More specifically, we set $\mathcal{D}_{l,t} = \tilde{S}_{l,t}/n_{l,t}$, where $\tilde{S}_{l,t} = B(\{x_j : i_j = l \wedge j \in [t]\})$ is the noisy sum output by the hybrid mechanism (which guarantees ϵ -differential privacy) and $n_{l,t}$ is the total number of times that s_l has been posted until time t .

According the above discussions, a straightforward idea is to directly select a price $s_{l'} \in \mathcal{S}$ at any time $t + 1$ such that $\min\{m\mathcal{D}_{l',t}, \frac{W}{s_{l'}}\}$ is maximized. However, this method cannot lead to good performance ratios, as the estimation $\mathcal{D}_{l,t}$ is not accurate for any l and t . Therefore, we introduce two factors $\mathcal{H}_{l,t} \triangleq \frac{\sqrt{8} \log 4t^4}{\epsilon n_{l,t}} (1 + \log n_{l,t})$ and $\sigma_{l,t} \triangleq \sqrt{\frac{5 \ln t}{2n_{l,t}}}$ as the “exploration factors” to compensate for the estimation error of $\mathcal{D}_{l,t}$, and hence select a price l to maximize

$$\tilde{\varphi}_{l,t} \triangleq \min\{m(\mathcal{D}_{l,t} + \sigma_{l,t} + \mathcal{H}_{l,t}), W/s_l\} \quad (5)$$

at any time $t + 1$. Indeed, the exploration factors $\mathcal{H}_{l,t}$ and $\sigma_{l,t}$ are elaborately designed to get a low “regret” of DPP-UCB, which will be seen from our performance analysis in Sec. 4.3.

With the intuitive ideas explained above, we describe the details of Algorithm 3 as follows. Algorithm 3 first tries each price in $\mathcal{S}$ to get the initial knowledge (line 4), and then selects $s_{i_t} \in \mathcal{S}$ as the posted price p_t for any user $t \in [m]$ according to Eqn. (5) (line 5). After observing any user t ’s response to the posted price, DPP-UCB

Algorithm 3: The DPP-UCB Algorithm

Input: $W, m, \mathcal{S}, \epsilon \in [0, 1]$
Output: $\mathbf{p}, \mathcal{N}, T$

```

1  $t \leftarrow 1, W_0 \leftarrow W, \forall i \in \{1, \dots, |\mathcal{S}|\} : n_{i,0} \leftarrow 0;$ 
2  $\mathbf{p} \leftarrow \mathbf{0}; \mathcal{N} \leftarrow \emptyset; T \leftarrow 0; \mathbf{x} \leftarrow \mathbf{0}$ 
3 while  $t \leq m$  do
4   if  $t \leq |\mathcal{S}|$  then  $i_t \leftarrow t;$ 
5   else  $i_t \leftarrow \arg \max_{l \in \{1, \dots, |\mathcal{S}|\}} \tilde{\varphi}_{l,t-1};$ 
6    $p_t \leftarrow s_{i_t};$ 
7   if  $s_{i_t} > W_{t-1}$  then break;
8   else  $T \leftarrow t;$ 
9   Post  $s_{i_t}$  to user  $t$  and update  $x_t;$ 
10  if  $x_t = 1$  then  $\mathcal{N} \leftarrow \mathcal{N} \cup \{t\};$ 
11   $W_t \leftarrow W_{t-1} - s_{i_t} \cdot x_t$  /* $W_t$ : leftover budget*/
12   $n_{i_t,t} \leftarrow n_{i_t,t-1} + 1; t \leftarrow t + 1$ 
13 return  $\mathbf{p}, \mathcal{N}, T$ 
```

updates the parameters (lines 9-12) and stops at a finite time $T \leq m$, such that user T is the last paid user.

Note that Algorithm 3 always selects each price in $\mathcal{S}$ sequentially in the initialization phase (i.e., when $t \leq |\mathcal{S}|$). Therefore, Algorithm 3 trivially achieves ϵ -differential privacy in the initialization phase due to the reason that the price selection therein is not affected by the users’ actions at all. After the initialization phase, Algorithm 3 selects prices based on the output of the hybrid mechanism, which achieves differential privacy. As DP is immune to post-processing [15], we can prove:

Theorem 6. DPP-UCB achieves ϵ -differential privacy.

Note that DPP-UCB decides the price p_t for any user t even before user t arrives. Thus, it can be easily proven that DPP-UCB achieves dominant strategy truthfulness and individual rationality as p_t is independent of user t . More specifically, any user $t \in [m]$ maximizes her/his utility u_t by truthfully reporting $x_t = \mathbf{1}(c_t \leq p_t)$ and then gets a non-negative utility $x_t(p_t - c_t)$, where $\mathbf{1}(\cdot)$ is the indicator function.

4.3 Performance Analysis on the Revenue

In this section, we show that DPP-UCB has a low “regret” on revenue, where the regret is defined as $\text{Reg} = \varphi_{l^*} - \mathbb{E}\{|\mathcal{N}|\}$ (recall that $\mathcal{N}$ is the winner set output by DPP-UCB). Intuitively, the regret measures how much revenue that we can lose compared with an omniscient algorithm that does not guarantee differential privacy. Before we describe our regret analysis, we quote the celebrated Hoeffding’s inequality in Lemma 2, which is used in our proofs:

Lemma 2 (Hoeffding’s inequality [29]). Let $Y_1, Y_2, \dots, Y_n$ be a sequence of random variables with common support $[0, 1]$. If $\mathbb{E}\{Y_i | Y_1, Y_2, \dots, Y_{i-1}\} \leq \Phi$ for any $i \leq n$, then we have $\Pr\{\frac{1}{n} \sum_{i=1}^n Y_i - \Phi \geq \kappa\} \leq \exp\{-2\kappa^2 n\}$ for any $\kappa > 0$. If $\mathbb{E}\{Y_i | Y_1, Y_2, \dots, Y_{i-1}\} \geq \Phi$ for any $i \leq n$, then we have $\Pr\{\frac{1}{n} \sum_{i=1}^n Y_i - \Phi \leq \kappa\} \leq \exp\{-2\kappa^2 n\}$ for any $\kappa > 0$.

The main idea of our regret analysis is that we first bound the expected number of the sub-optimal prices selected from $\mathcal{S} \setminus \{s_{l^*}\}$ by the DPP-UCB algorithm and then bound the revenue loss suffered from selecting these sub-optimal prices. However, our analysis is very different from that in the traditional MAB problem [19], [25], as we need to take DP into account and handle the sub-optimal prices in $\{s_l \in \mathcal{S} | l > l^*\}$ and $\{s_l \in \mathcal{S} | l < l^*\}$ separately. In the sequel, we give the regret analysis in detail.

We first give some bounds on the probabilities of the suboptimal prices (i.e., prices in $\mathcal{S} \setminus \{s_{l^*}\}$) being selected in DPP-UCB, as shown by Lemma 5-6. Intuitively, Lemma 5 and Lemma 6 imply that, after the suboptimal prices in $\mathcal{S} \setminus \{l^*\}$ have been selected a sufficient number of times in the history, the probability of selecting them again by DPP-UCB is very small, as DPP-UCB has acquired “sufficiently accurate” knowledge about these prices to determine that they are sub-optimal. The proof of Lemma 5 leverages Lemma 4, which uses a property of the celebrated Lambert function stated in Lemma 3:

Lemma 3. [30], [31] Let $\bar{W}$ be the Lambert function which satisfies $\forall x \in \mathbb{R} : x = \bar{W}(xe^x)$. The equation $e^{-cx} = a_0(x-r)$ ($a_0, c, r \in \mathbb{R}; a_0 \neq 0$) has the solution $x = r + \frac{1}{c}\bar{W}\left(\frac{ce^{-cr}}{a_0}\right)$.

Lemma 4. Define $\Delta_l = \varphi_{l^*} - \varphi_l$, $Q_l(t) = \frac{2\sqrt{8} \ln 4t^4}{\epsilon(1-\nu)(\Delta_l/m)}$ and $\zeta(l, t) = \max\left(\frac{10 \ln t}{\nu^2(\Delta_l/m)^2}, Q_l(t)(\ln Q_l(t) + 7)\right)$ where ν is any number in $(0, 1)$. When $n_{l,t} > \zeta(l, t)$, we must have $\Delta_l/m > 2\sigma_{l,t} + 2\mathcal{H}_{l,t}$

Proof: Based on the properties of the celebrated Lambert function [32] and Lemma 3, we know that a sufficient condition for the inequality $i > y \ln i + y$ ($\forall i \in \mathbb{N}^+, \forall y > 0$) to hold is $i > y(\ln y + 7)$. Therefore, if $n_{l,t} > Q_l(t)(\ln Q_l(t) + 7)$, we must have $n_{l,t} > Q_{l,t}(1 + \ln n_{l,t})$, which implies $(1-\nu)\Delta_l/m > 2\mathcal{H}_{l,t}$. In addition, it can be easily known that $\nu\Delta_l/m > 2\sigma_{l,t}$ when $n_{l,t} > \frac{10 \ln t}{\nu^2(\Delta_l/m)^2}$. Hence, the lemma follows. $\square$

Lemma 5. For any $t \geq k$ and any $j < l^*$, we have $\mathbb{P}\{i_{t+1} = j; n_{j,t} \geq \zeta(j, t)\} \leq 6t^{-4}$.

Lemma 6. For any $t \geq k$ and any $j > l^*$, we have $\mathbb{P}\{i_{t+1} = j\} \leq 2t^{-4}$.

With Lemma 4-6, we are able to bound the expectations of the numbers of the sub-optimal prices selected by DPP-UCB, which is shown in Lemma 8-9. Intuitively, Lemma 8 shows that the expected number of times that any sub-optimal price $j < l^*$ is selected in DPP-UCB is bounded by $\mathcal{O}(\zeta(j, \beta))$, while Lemma 9 shows that the expected number of times that any sub-optimal price $j > l^*$ is selected in DPP-UCB is bounded by a constant. The proof of Lemma 8 leverages Lemma 7, which essentially provides an upper bound on DPP-UCB’s stopping time (i.e., T).

Lemma 7. Let $\beta = \lceil 2W/(s_1 D_1) \rceil$. Then, we have

$$\sum_{t=\beta}^{\infty} \mathbb{P}\{W_t \geq 0\} \leq 2s_k^2/(s_1 D_1)^2$$

Lemma 8. For any $j < l^*$, we have $\mathbb{E}\{n_{j,T}\} \leq 2\lceil \zeta(j, \beta) \rceil + \frac{2s_k^2}{(s_1 D_1)^2} + \frac{\pi^4}{15}$.

Proof: For any $j < l^*$, we have

$$\begin{aligned} n_{j,\beta} &= 1 + \sum_{t=k}^{\beta-1} \mathbf{1}\{i_{t+1} = j\} \\ &= 1 + \sum_{t=k}^{\beta-1} \mathbf{1}\{i_{t+1} = j; n_{j,\beta} < \zeta(j, \beta)\} \\ &\quad + \sum_{t=k}^{\beta-1} \mathbf{1}\{i_{t+1} = j; n_{j,\beta} \geq \zeta(j, \beta)\} \\ &\leq \lceil \zeta(j, \beta) \rceil + \sum_{t=k}^{\beta-1} \mathbf{1}\{i_{t+1} = j; n_{j,\beta} \geq \zeta(j, \beta)\} \end{aligned} \quad (6)$$

If $n_{j,\beta} \geq \zeta(j, \beta)$, there must exist certain $v \in [k, \beta - 1]$ such that $n_{j,v} < \zeta(j, \beta)$ and $n_{j,v+1} \geq \zeta(j, \beta)$. Therefore, we have

$$\begin{aligned} &\sum_{t=k}^{\beta-1} \mathbf{1}\{i_{t+1} = j; n_{j,\beta} \geq \zeta(j, \beta)\} \\ &\leq \lceil \zeta(j, \beta) \rceil + \sum_{t=v+1}^{\beta-1} \mathbf{1}\{i_{t+1} = j; n_{j,\beta} \geq \zeta(j, \beta)\} \\ &\leq \lceil \zeta(j, \beta) \rceil + \sum_{t=v+1}^{\beta-1} \mathbf{1}\{i_{t+1} = j; n_{j,t} \geq \zeta(j, \beta)\} \\ &\leq \lceil \zeta(j, \beta) \rceil + \sum_{t=k}^{\beta-1} \mathbf{1}\{i_{t+1} = j; n_{j,t} \geq \zeta(j, t)\} \end{aligned} \quad (7)$$

Combining equation (6)-(7) with Lemma 5 and Lemma 7 gives us

$$\begin{aligned} &\mathbb{E}\{n_{j,T}\} \\ &\leq \mathbb{E}\{n_{j,\beta}\} + \mathbb{E}\left\{\sum_{t=\beta}^{\infty} \mathbf{1}\{i_t = j\}\right\} \\ &\leq 2\lceil \zeta(j, \beta) \rceil + \mathbb{E}\left\{\sum_{t=k}^{\beta-1} \mathbf{1}\{i_{t+1} = j; n_{j,t} \geq \zeta(j, t)\}\right\} \\ &\quad + \mathbb{E}\left\{\sum_{t=\beta}^{\infty} \mathbb{P}\{W_t \geq 0\}\right\} \\ &\leq 2\lceil \zeta(j, \beta) \rceil + 2s_k^2/(s_1 D_1)^2 + \sum_{t=1}^{\infty} 6t^{-4} \\ &\leq 2\lceil \zeta(j, \beta) \rceil + 2s_k^2/(s_1 D_1)^2 + \frac{\pi^4}{15} \end{aligned} \quad (8)$$

where (8) is due to the Riemann zeta function $\sum_{t=1}^{\infty} t^{-4} = \frac{\pi^4}{90}$ [19]. Hence, the lemma follows. $\square$

Lemma 9. For any $j > l^*$, we have $\mathbb{E}\{n_{j,T}\} \leq 1 + \frac{\pi^4}{45}$

Proof: Using Lemma 6 and the Riemann zeta function, we can get $\mathbb{E}\{n_{j,T}\} = \mathbb{E}\{1 + \sum_{t=k}^{T-1} \mathbf{1}\{i_{t+1} = j\}\} \leq 1 + \sum_{t=1}^{\infty} 2t^{-4} = 1 + \frac{\pi^4}{45}$. $\square$

Now, we are ready to give the regret bound of DPP-UCB, as shown by Theorem 7. The main idea for proving Theorem 7 is that, as selecting l^* incurs zero regret, we only need to consider the regret caused by selecting the sub-optimal prices in $\mathcal{S} \setminus \{l^*\}$, which can be calculated by using Lemma 8 and Lemma 9. To derive the regret bound shown in Theorem 7, we also need a lower bound on the stopping time T of DPP-UCB, which is presented in Lemma 10:

Lemma 10. *The stopping time T of the DPP-UCB algorithm is bounded by*

$$\mathbb{E}\{T\} > \min\{W/(D_{l^*} s_{l^*}), m\} - \theta_2/(D_{l^*} s_{l^*}) - \sum_{j>l^*} \mathbb{E}\{n_{j,T}\} \left(\frac{D_j s_j}{D_{l^*} s_{l^*}} - 1 \right)$$

Theorem 7. *The regret of the DPP-UCB algorithm has an upper bound of $\mathcal{O}(\log W \log \log W)$.*

Proof: Note that $|\mathcal{N}| = \sum_{t=1}^T x_t$. Therefore, we have

$$\begin{aligned} \text{Reg} &\leq \varphi(s_{l^*}) - \mathbb{E}\left\{\sum_{t=1}^T x_t\right\} \\ &= \min\{mD_{l^*}, W/s_{l^*}\} - \mathbb{E}\{T\}D_{l^*} \\ &\quad + \sum_{j \in [k]} \mathbb{E}\{n_{j,T}\} (D_{l^*} - D_j) \\ &\leq \min\{mD_{l^*}, W/s_{l^*}\} - \mathbb{E}\{T\}D_{l^*} \\ &\quad + \sum_{j < l^*} \mathbb{E}\{n_{j,T}\} (D_{l^*} - D_j) \\ &\leq \frac{\theta_2}{s_{l^*}} + \sum_{j < l^*} \mathbb{E}\{n_{j,T}\} (D_{l^*} - D_j) \\ &\quad + \frac{1}{s_{l^*}} \sum_{j > l^*} \mathbb{E}\{n_{j,T}\} (D_j s_j - D_{l^*} s_{l^*}) \quad (9) \\ &\leq \frac{\theta_2}{s_{l^*}} + \sum_{j < l^*} (D_{l^*} - D_j) \left(2[\zeta(j, \beta)] + \frac{2s_k^2}{(s_1 D_1)^2} \right. \\ &\quad \left. + \frac{\pi^4}{15} \right) + \sum_{j > l^*} \left(\frac{D_j s_j}{s_{l^*}} - D_{l^*} \right) \left(1 + \frac{\pi^4}{45} \right) \quad (10) \end{aligned}$$

where (9) is due to Lemma 10 and (10) is due to Lemma 8-9. Note that the factor $\zeta(j, \beta)$ is in the order of $\mathcal{O}(\log W \log \log W)$. Hence, the theorem follows. $\square$

It can be seen from Theorem 7 that the average regret of DPP-UCB asymptotically approaches zero; i.e., we have $\lim_{W \rightarrow \infty} \text{Reg}/W = 0$. This suggests that DPP-UCB is a Hannan-consistent learning algorithm [19].

Remark: Until now, we have assumed that the number of users m is fixed and known under the online setting. Although this assumption has also been widely adopted in the literature (e.g., [20], [28], [33], [34]), it may not hold in some crowdsourcing applications where the number of participants is stochastic. Fortunately, in such cases, we can often get the distribution knowledge and hence the expected value of the number of participants. For example, the historical mobility traces of the users could be used to estimate the number of users appeared in the Points of Interests (PoI) of mobile crowdsourcing applications [9], [35].

Based on the above observation, we can use a simple method to extend our DPP-UCB algorithm to the case that the number of users is stochastic with known distributions. More specifically, we can simply replace m by its expectation in our algorithms, while all our performance analysis and performance bounds remain the same. This is due to the reason that, the revenue of any pricing strategy under the online setting is originally defined as an expected value (see Sec. 2.2), so we should replace m by its expected value for calculating the revenue of any price in $\mathcal{S}$ (including the optimal price) when m is stochastic. As such, the expected value of m can be

Fig. 3: Distribution of the traveling distances in the T-Drive dataset (kilometers)

considered as a constant in our algorithms, so all our performance analysis still holds under this setting.

5 PERFORMANCE EVALUATION

We conduct extensive experiments to study the performance of our mechanisms. The purpose of our experiments is to compare our algorithms with the related work on revenue and regret, using both synthetic datasets and real datasets.

5.1 Experimental Settings

In the experiments, we compare OPEX with R_{opt} and PWDP in the offline setting, and compare DPP-UCB with several representative algorithms proposed in the literature including BP-UCB [20], UCB-BV2 [36] and MRCB [37]. BP-UCB is a dynamic pricing algorithm for customers arriving online, while UCB-BV2 and MRCB are two representative budgeted MAB algorithms with random arm costs. As neither of BP-UCB, MRCB and UCB-BV2 guarantees differential privacy, we have to adapt them to achieve DP for fair comparison. Therefore, we incorporate the hybrid mechanism into BP-UCB, MRCB and UCB-BV2, i.e., using the noisy means as the estimations of the expectations of the arms' costs (as we do in DPP-UCB). We also implement OPT^* , which is the benchmark algorithm that always selects the price s_{l^*} in the online setting.

We use both a real-world dataset and synthetic datasets to test the performance of our implemented algorithms. The real-world dataset used by us is T-Drive [38], which is a mobile trajectory dataset published by Microsoft Research. This dataset contains the GPS trajectories of 10,357 taxis in the Beijing city. In our experiments, we consider the drivers in the T-Drive dataset as mobile crowdsourcing users, and their costs are set in proportional to their travelling distances. The distribution of the users' travelling distances in the T-Drive dataset is plotted in Fig. 3.

In each of the generated synthetic datasets, the cost of any user is sampled from one of the following distributions (by a uniformly random selection): the Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ (truncated to have the support $[0, 1]$), the uniform distribution $\mathcal{U}(0, 1)$ and the beta distribution $\text{Beta}(\alpha, \beta)$. Both of the parameters μ, σ are randomly sampled from the uniform distribution $\mathcal{U}(0, 1)$, while α, β are randomly sampled from $\mathcal{U}(0, 10)$. To achieve unbiased performance comparison, the reported data for the synthetic datasets are the average

Fig. 4: Comparing the revenue of the implemented offline algorithms using the T-Drive dataset

Fig. 5: Comparing the revenue/regret of the implemented online algorithms using the T-Drive dataset

running results on 200 datasets. We also follow [20] to set $\theta_1 = 0.01$, $\theta_2 = 1$ and $k = 20$ for all the implemented algorithms.

5.2 Experimental Results

5.2.1 Overall Performance

In Fig. 4, we compare the performance of the implemented algorithms under the offline setting using the T-Drive dataset, where the budget W is scaled from 100 to 1000 and ϵ is set to 1, 0.6 and 0.2 in Fig 4(a)-4(c), respectively. It can be seen that the revenue of all offline algorithms increases with W , which can be explained by the reason that more crowdsourcing users can be recruited when W increases. Another fact revealed by Fig. 4 is that the revenue of OPEX does not vary much when ϵ decreases, which demonstrates the robustness of OPEX against the variation of ϵ . Finally, it can be seen from Fig. 4 that the performance of OPEX is very close to PWDP, which indicates that OPEX can achieve differential privacy without losing too much revenue.

In Fig. 5, we study the revenue and regret performance of the implemented algorithms under the online setting using the T-Drive dataset, where the parameter settings are the same to those in Fig. 4. It can be seen from Fig. 5 that the revenue of all algorithms increases with the budget, while DPP-UCB outperforms BP-UCB, UCB-BV2 and MRCB. We also notice that a smaller ϵ results in smaller revenue of all algorithms except for OPT*, which is not surprising as the algorithms need to spend more

budget to identify the optimal price when ϵ decreases. Finally, the results shown in Fig. 5 demonstrate that the average regret of DPP-UCB (i.e., Reg/W) approaches 0 when the budget increases, while it is much smaller than the average regret of the other two algorithms. This corroborates the logarithmic regret bound of DPP-UCB proved in Sec. 4.3.

In Figs. 6-7, we study the performance of the implemented algorithms using the synthetic datasets. Following the work in [20], the number of users in Figs. 6-7 is set to W/θ_1 , and the other parameter settings are the same to those in Figs. 4-5. It can be seen that the results shown in Figs. 6-7 are similar to those in Figs. 4-5, and our algorithms outperform the baseline algorithms both under the online setting and under the offline setting. This can be explained by similar reasons with those for Figs. 4-5.

5.2.2 Privacy Leakage

In this section, we study how the users' privacy is protected by the implemented algorithms. Following [8], we use the Kullback-Leibler (KL) divergence to measure the Privacy Leakage (PL) of the algorithms. More specifically, we define PL under the offline setting as follows:

Definition 5. (Privacy leakage under the offline setting) Under the offline setting, the privacy leakage with respect to any two adjacent bidding vectors $\mathbf{b}$ and $\mathbf{b}'$ (i.e., $\mathbf{b}$ and $\mathbf{b}'$ differ in only one user's bid) is defined as $PL =$

Fig. 6: Comparing the revenue of the implemented offline algorithms using the synthetic datasets

Fig. 7: Comparing the revenue/regret of the implemented online algorithms using the synthetic datasets

$$\sum_{\mathbf{p} \in \mathcal{S}^m} \mathbb{P}(\mathbf{p}|\mathbf{b}) \ln \left[\frac{\mathbb{P}(\mathbf{p}|\mathbf{b})}{\mathbb{P}(\mathbf{p}|\mathbf{b}')} \right]$$

Intuitively, the privacy leakage defined above measures how much the payment profile varies when one user's bid changes. Similarly, we define PL under the online setting as follows:

Definition 6. (Privacy leakage for the online setting) Under the online setting, the privacy leakage with respect to any $\mathbf{p}_{j-1}$ and any adjacent pair $(\mathbf{x}_{j-1}, \mathbf{x}'_{j-1})$ (see Definition 4) is defined as $PL = \sum_{p_j \in \mathcal{S}} \mathbb{P}(p_j|\mathbf{p}_{j-1}, \mathbf{x}_{j-1}) \ln \left[\frac{\mathbb{P}(p_j|\mathbf{p}_{j-1}, \mathbf{x}_{j-1})}{\mathbb{P}(p_j|\mathbf{p}_{j-1}, \mathbf{x}'_{j-1})} \right]$

Note that PWDP is a deterministic algorithm and has infinite PL. Therefore, we only compare the PL of the implemented randomized algorithms in Fig. 8, where the parameter settings are the same to those in Figs. 4-7. To achieve unbiased comparison, we randomly generate 1000 adjacent bidding vectors (or 1000 adjacent pairs under the online setting), and plot the average PL in Fig. 8. The results in Fig. 8 reveal that, the privacy leakage of all the implemented algorithms increases with ϵ , as a larger ϵ allows for more privacy leakage according to the definition of DP [15]. However, the privacy leakage of our algorithms (i.e., OPEX and DPP-UCB) is significantly smaller than those of the other algorithms, which demonstrates the superiority of our approach for privacy preservation.

Fig. 8: Comparing the privacy leakage

6 RELATED WORK

The incentivization problem in mobile crowdsourcing has been extensively studied, and the related studies in this area can be found in two excellent surveys [39], [40]. In particular, the insightful survey in [39] has presented a novel and comprehensive taxonomy of existing incentive mechanisms for mobile crowdsourcing systems, and it has also discussed about the related approaches in depth. The more recent survey in [40] has proposed the first framework in the literature for defining and enforcing Quality of Information (QoI) in mobile crowdsourcing, and has also proposed some novel research challenges and possible research directions in this area.

Among the existing crowdsourcing incentivization approaches, the work in [6], [7], [16], [17] has designed truthful auction mechanisms under a limited budget. Most of these studies are based on the framework of budget feasible mechanisms proposed in [2]–[4]. It can be

seen that the best approximation ratio of the budget feasible mechanisms investigated in [2]–[4] is $\frac{1}{3}$ (presented in [3]). However, all these related work has neglected the privacy issue.

Recently, the privacy protection problems in mobile crowdsourcing have begun to arouse interests in the literature. Two excellent studies [41], [42] have provided insightful surveys on the methods and challenges for protecting the privacy of mobile crowdsourcing users. The seminal work in [43] has proposed FIDES, a brilliant trust-based framework for secure user incentivization in mobile crowdsourcing, which is the first trust-based framework that simultaneously solves the problems of incentivizing users' participation and guaranteeing data reliability. The work in [44] has proposed a smart anonymity-preserving reputation framework for mobile crowdsourcing, which is agnostic to both the reputation assignment algorithm and the crowdsourcing application. The work in [45], [46] has studied the privacy preserving data publishing and aggregation problems in mobile crowdsourcing, but without considering the incentivization problem. The excellent studies in [35], [47] have proposed some novel approaches to protect the location privacy of the users, using the tools of differential privacy or k -anonymity. Two closest studies to ours are [8] and [18], where some ingenious auction mechanisms are proposed to protect the bidding privacy of the users in mobile crowdsourcing. However, no budget constraints are considered in [8], [18], and their problem models and optimization goals are both very different from ours. Indeed, both [8] and [18] aim to minimize the total payment/cost in mobile crowdsourcing, while our goal is to maximize the system revenue under a budget constraint. Moreover, only offline algorithms are proposed in [8], [18]. Due to these essential discrepancies, the algorithms and methods proposed in [8], [18] cannot be applied to our case.

The dynamic pricing problem for online customers has also been studied in the literature [20], [28], [33], [34]. Nevertheless, these studies neglected either the budget constraint or the privacy issue. Meanwhile, it is noted that the dynamic pricing problem can be considered as a variant of the budgeted multi-armed bandit problem investigated in [36], [48], [49]. Nevertheless, none of these studies considered the privacy issue. Moreover, our online pricing problem has some unique features ignored by [36], [48], [49]; i.e., the rewards for selecting different prices/arms are implicitly correlated (as any user rejecting a given price would also reject a lower price). Therefore, our regret analysis for DPP-UCB is very different from those in [36], [48], [49].

Finally, it can be seen that most of the studies on multi-armed bandits have not considered the privacy issue with only a few exceptions [27], [50]. However, both [50] and [27] assume an unlimited budget for playing the arms, so their algorithms cannot be used to address our problem.

7 CONCLUSION AND DISCUSSIONS

We have studied the problem of designing differentially private incentivization mechanisms for mobile crowdsourcing, where the total payment to users should not exceed a predefined budget. We have proposed novel algorithms for our problem both in an offline setting and an online setting. We have shown that our mechanisms achieve provable theoretical performance bounds on revenue, truthfulness and differential privacy simultaneously, and the effectiveness of our approach has also been corroborated by the results of numerical experiments. To the best of our knowledge, we are the first to propose privacy-preserving mechanisms with provable performance bounds for budget limited mobile crowdsourcing.

Although the effectiveness of our algorithms has been proved by both theoretical analysis and experimental evaluations, there are still several improvements that could be done. First, as our model only assumes homogeneous crowdsourcing tasks (and hence homogeneous task values), extending the current model to the case of heterogeneous tasks would be an interesting problem. Second, it would be interesting to design more efficient algorithms for our problem under the continuous-pricing scenario, as the running time of the current algorithms could be high in this case. Third, although we have followed some related work (e.g., [20], [28], [33], [34]) to assume that there are some prior knowledge on the number of users under the online setting, designing algorithms for our problem without this assumption could make our approach more general. We plan to investigate all these issues in the future.

ACKNOWLEDGMENTS

This work is partially supported by National Natural Science Foundation of China under Grant No. 61772491, No. 61472460, and Natural Science Foundation of Jiangsu Province under Grant No. BK20161256.

REFERENCES

- [1] B. Guo, Z. Wang, Z. Yu, Y. Wang, N. Y. Yen, R. Huang, and X. Zhou, "Mobile crowd sensing and computing: The review of an emerging human-powered sensing paradigm," *ACM Computing Surveys*, vol. 48, no. 1, pp. 7:1–7:31, 2015.
- [2] Y. Singer, "Budget feasible mechanisms," in *FOCS*, 2010, pp. 765–774.
- [3] N. Chen, N. Gravin, and P. Lu, "On the Approximability of Budget Feasible Mechanisms," in *SODA*, 2011, pp. 685–699.
- [4] X. Bei, N. Chen, N. Gravin, and P. Lu, "Budget feasible mechanism design: from prior-free to bayesian," in *STOC*, 2012, pp. 449–458.
- [5] T. Luo, S. S. Kanhere, H. P. Tan, F. Wu, and H. Wu, "Crowdsourcing with tullock contests: A new perspective," in *INFOCOM*, 2015, pp. 2515–2523.
- [6] D. Zhao, X. Y. Li, and H. Ma, "Budget-feasible online incentive mechanisms for crowdsourcing tasks truthfully," *IEEE/ACM Transactions on Networking*, vol. 24, no. 2, pp. 647–661, 2016.
- [7] D. Yang, G. Xue, X. Fang, and J. Tang, "Crowdsourcing to Smartphones: Incentive Mechanism Design for Mobile Phone Sensing," in *MobiCom*, 2012, pp. 173–184.
- [8] H. Jin, L. Su, B. Ding, K. Nahrstedt, and N. Borisov, "Enabling privacy-preserving incentives for mobile crowd sensing systems," in *ICDCS*, 2016, pp. 344–353.

- [9] P. Cheng, X. Lian, Z. Chen, R. Fu, L. Chen, J. Han, and J. Zhao, "Reliable diversity-based spatial crowdsourcing by moving workers," in *Vldb*, 2015, pp. 1022–1033.
- [10] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec, "Protecting location privacy: optimal strategy against localization attacks," in *CCS*, 2012, pp. 617–627.
- [11] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*, 2007.
- [12] L. Sweeney, "k-anonymity: A model for protecting privacy," *World Scientific International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.
- [13] A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkatasubramanian, "l-diversity: Privacy beyond k-anonymity," in *ICDE*, 2006, pp. 24–24.
- [14] C. Clifton and T. Tassa, "On syntactic anonymity and differential privacy," *Transactions on Data Privacy*, vol. 6, no. 2, pp. 161–183, 2013.
- [15] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [16] Q. Zhang, Y. Wen, X. Tian, X. Gan, and X. Wang, "Incentivize crowd labeling under budget constraint," in *INFOCOM*, 2015, pp. 2812–2820.
- [17] Y. Singer and M. Mittal, "Pricing mechanisms for crowdsourcing markets," in *WWW*, 2013, pp. 1157–1166.
- [18] J. Lin, D. Yang, M. Li, J. Xu, and G. Xue, "Bidguard: A framework for privacy-preserving crowdsensing incentive mechanisms," in *CNS*, 2016, pp. 145–153.
- [19] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *Foundations and Trends® in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.
- [20] A. Singla and A. Krause, "Truthful incentives in crowdsourcing tasks using regret minimization mechanisms," in *WWW*, 2013, pp. 1167–1178.
- [21] <https://www.mturk.com/>.
- [22] Y. Gao and A. Parameswaran, "Finish them!: Pricing algorithms for human computation," *Proceedings of the VLDB Endowment*, vol. 7, no. 14, pp. 1965–1976, 2014.
- [23] <https://www.google.com/recaptcha/intro/>.
- [24] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *FOCS*, 2007, pp. 94–103.
- [25] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2–3, pp. 235–256, 2002.
- [26] T.-H. H. Chan, E. Shi, and D. Song, "Private and continual release of statistics," *ACM Transactions on Information and System Security*, vol. 14, no. 3, p. 26, 2011.
- [27] A. C. Y. Tossou and C. Dimitrakakis, "Algorithms for differentially private multi-armed bandits," in *AAAI*, 2016, pp. 2087–2093.
- [28] A. Badanidiyuru, R. Kleinberg, and Y. Singer, "Learning on a Budget: Posted Price Mechanisms for Online Procurement," in *EC*, 2012, pp. 128–145.
- [29] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13–30, 1963.
- [30] D. Barry, J.-Y. Parlange, L. Li, H. Prommer, C. Cunningham, and F. Stagnitti, "Analytical approximations for real values of the lambert w-function," *Mathematics and Computers in Simulation*, vol. 53, no. 1, pp. 95–103, 2000.
- [31] S. Yi, P. W. Nelson, and A. G. Ulsoy, *Time-delay systems: analysis and control using the Lambert W function*. World Scientific, 2010.
- [32] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth, "On the lambertW function," *Advances in Computational Mathematics*, vol. 5, no. 1, pp. 329–359, 1996.
- [33] M. Babaioff, S. Dughmi, R. Kleinberg, and A. Slivkins, "Dynamic pricing with limited supply," *ACM Transactions on Economics and Computation*, vol. 3, no. 1, p. 4, 2015.
- [34] R. Kleinberg and T. Leighton, "The value of knowing a demand curve: Bounds on regret for online posted-price auctions," in *FOCS*, 2003, pp. 594–605.
- [35] X. Jin and Y. Zhang, "Privacy-preserving crowdsourced spectrum sensing," in *INFOCOM*, 2016, pp. 2322–2330.
- [36] W. Ding, T. Qin, X. Zhang, and T. Liu, "Multi-armed bandit with budget constraint and variable costs," in *AAAI*, 2013, pp. 232–238.
- [37] Y. Xia, T. Qin, W. Ma, N. Yu, and T. Liu, "Budgeted multi-armed bandits with multiple plays," in *IJCAI*, 2016, pp. 2210–2216.
- [38] <https://www.microsoft.com/en-us/research/publication/t-drive-trajectory-data-sample/>.
- [39] F. Restuccia, S. K. Das, and J. Payton, "Incentive mechanisms for participatory sensing: Survey and research challenges," *ACM Transactions on Sensor Networks*, vol. 12, no. 2, pp. 13:1–13:40, 2016.
- [40] F. Restuccia, N. Ghosh, S. Bhattacharjee, S. K. Das, and T. Melodia, "Quality of information in mobile crowdsensing: Survey and research challenges," *ACM Transactions on Sensor Networks*, vol. 13, no. 4, pp. 34:1–34:43, 2017.
- [41] D. Christin, "Privacy in mobile participatory sensing: Current trends and future challenges," *Elsevier Journal of Systems and Software*, vol. 116, pp. 57–68, 2016.
- [42] L. Pournajaf, D. A. Garcia-Ulloa, L. Xiong, and V. Sunderam, "Participant privacy in mobile crowd sensing task management: A survey of methods and challenges," *ACM SIGMOD Record*, vol. 44, no. 4, pp. 23–34, 2016.
- [43] F. Restuccia and S. K. Das, "Fides: A trust-based framework for secure user incentivization in participatory sensing," in *WoWMoM*, 2014, pp. 1–10.
- [44] D. Christin, C. Roßkopf, M. Hollick, L. A. Martucci, and S. S. Kanhere, "Incognisense: An anonymity-preserving reputation framework for participatory sensing applications," *Elsevier Pervasive and mobile Computing*, vol. 9, no. 3, pp. 353–371, 2013.
- [45] Q. Wang, Y. Zhang, X. Lu, Z. Wang, Z. Qin, and K. Ren, "Rescuedp: Real-time spatio-temporal crowd-sourced data publishing with differential privacy," in *INFOCOM*, 2016, pp. 1–9.
- [46] G. Zhuo, Q. Jia, L. Guo, M. Li, and P. Li, "Privacy-preserving verifiable data aggregation and analysis for cloud-assisted mobile crowdsourcing," in *INFOCOM*, 2016, pp. 1953–1961.
- [47] D. Yang, X. Fang, and G. Xue, "Truthful incentive mechanisms for k-anonymity location privacy," in *INFOCOM*, 2013, pp. 2994–3002.
- [48] L. Tran-Thanh, A. Chapman, E. M. de Cote, A. Rogers, and N. R. Jennings, "Epsilon-first policies for budget-limited multi-armed bandits," in *AAAI*, 2010, pp. 1211–1216.
- [49] L. Tran-Thanh, A. Chapman, A. Rogers, and N. R. Jennings, "Knapsack based optimal policies for budget-limited multi-armed bandits," in *AAAI*, 2012, pp. 1134–1140.
- [50] N. Mishra and A. Thakurta, "(nearly) optimal differentially private stochastic multi-arm bandits," in *UAI*, 2015, pp. 592–601.

Kai Han received his B.S. and Ph.D. degrees in computer science from the University of Science and Technology of China (USTC), Hefei, China, in 1997 and 2004, respectively. He is currently a Professor at the School of Computer Science and Technology, USTC. His research interests include wireless ad hoc and sensor networks, mobile and cloud computing, combinatorial and stochastic optimization, algorithmic game theory, and machine learning. He is a Member of both IEEE and ACM.

Huan Liu received his B.E. degree in Computer Science from Jilin University, China, in 2015. He is currently a master's student at the School of Computer Science and Technology, USTC. His research interests include mobile crowdsourcing and privacy-preserving data management.

Shaojie Tang is currently an assistant professor at the Naveen Jindal School of Management at the University of Texas at Dallas. He received his PhD in computer science from the Illinois Institute of Technology in 2012. His research interests include social networks, mobile commerce, game theory, e-business and optimization.

Mingjun Xiao is an associate professor at the School of Computer Science and Technology at the University of Science and Technology of China (USTC). He received his Ph.D. from USTC in 2004. His research interests include mobile computing and vehicular ad hoc networks.

Jun Luo received his PhD degree in computer science from EPFL, Lausanne, Switzerland in 2006. In 2008, he joined the faculty of the School of Computer Science and Engineering, Nanyang Technological University in Singapore, where he is currently an Associate Professor. His research interests include wireless networking and mobile computing, applied operations research and network security.