



Privacy-Preserving Batch-based Task Assignment in Spatial Crowdsourcing with Untrusted Server

Maocheng Li
csmichael@cse.ust.hk
HKUST
Hong Kong, China

Jiachuan Wang
jwangey@cse.ust.hk
HKUST
Hong Kong, China

Libin Zheng
zhenglb6@mail.sysu.edu.cn
Sun Yat-sen University
Guangzhou, China

Han Wu
han.wu@stu.ecnu.edu.cn
East China Normal
University
Shanghai, China

Peng Cheng
pcheng@sei.ecnu.edu.cn
East China Normal
University
Shanghai, China

Lei Chen
leichen@cse.ust.hk
HKUST
Hong Kong, China

Xuemin Lin
lxue@cse.unsw.edu.au
The University of New
South Wales
Sydney, Australia

ABSTRACT

In this paper, we study the privacy-preserving task assignment problem in spatial crowdsourcing, where the locations of both workers and tasks, prior to their release to the server, are perturbed with Geo-Indistinguishability (a differential privacy notion for location-based systems). Different from the previously studied online setting, where each task is assigned immediately upon arrival, we target the batch-based setting, where the server maximizes the number of successfully assigned tasks after a *batch* of tasks arrive. To achieve this goal, we propose the *k*-Switch solution, which first divides the workers into small groups based on the perturbed distance between workers/tasks, and then utilizes Homomorphic Encryption (HE) based secure computation to enhance the task assignment. Furthermore, we expedite HE-based computation by limiting the size of the small groups under *k*. Extensive experiments demonstrate that, in terms of the number of successfully assigned tasks, the *k*-Switch solution improves batch-based baselines by 5.9× and the existing online solution by 1.74×, with no privacy leak.

CCS CONCEPTS

• **Security and privacy** → **Domain-specific security and privacy architectures**; • **Information systems** → **Location based services**; **Crowdsourcing**.

KEYWORDS

Spatial crowdsourcing; differential privacy; task assignment

ACM Reference Format:

Maocheng Li, Jiachuan Wang, Libin Zheng, Han Wu, Peng Cheng, Lei Chen, and Xuemin Lin. 2021. Privacy-Preserving Batch-based Task Assignment in Spatial Crowdsourcing with Untrusted Server. In *Proceedings of the 30th*

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

CIKM '21, November 1–5, 2021, Virtual Event, Australia.

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8446-9/21/11...\$15.00

<https://doi.org/10.1145/3459637.3482288>

ACM Int'l Conf. on Information and Knowledge Management (CIKM '21), November 1–5, 2021, Virtual Event, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3459637.3482288>

1 INTRODUCTION

The mass adoption of GPS-equipped smart phones enables individuals to collaborate, participate, consume and produce valuable information about the environment and themselves. *Spatial crowdsourcing* (SC) systems (e.g., Foursquare [1], Gigwalk [2], MediaQ [13], and gMission [6]) have emerged to support such collaborations by assigning tasks to proper workers. Task assignment is a core issue in SC systems, asking the workers to move physically to specified locations to execute the tasks [12, 22].

To enable the SC server to properly assign tasks, in general, workers need to upload their locations. However, users' location is highly sensitive, because it can indicate users' whereabouts, and even disclose their private attributes. For example, visiting an urgent care center reveals certain medical conditions [9]. The server, which receives the locations, is untrusted and can be vulnerable to attacks. Thus, in previous studies [18, 21], privacy-preserving task assignment is proposed to enable users (workers/task requesters) to perturb their locations with *Geo-Indistinguishability* (Geo-I) [5] and upload only the perturbed locations. Geo-I is a widely adopted differential privacy notion [8] for location-based systems, which defends users' locations against strong adversaries with any prior knowledge [9].

In this paper, we study the task assignment problem with users' perturbed locations under the batch-based setting. That is, the SC server assigns the tasks batch-by-batch to the workers, with locations of both workers and tasks perturbed by Geo-I. Although Geo-I is an existing technique to protect location privacy of users, directly applying Geo-I to task assignment problem could lead to poor performance as measured by the number of successfully assigned tasks, because only the perturbed locations are available as the problem inputs. To the best of our knowledge, existing works on spatial crowdsourcing with Geo-I [18, 21] target the online setting where tasks come on the fly and are processed one by one. Batch-based task assignment with Geo-I-perturbed locations has not yet been studied. Directly applying the online assignment methods to the batch-based setting can result in poor solutions, because online

solution assigns tasks one by one, and previous assignments can not be changed when new tasks arrive. An example of the online solution making sub-optimal assignments is shown in Figure 1.




Figure 1: Online solution makes sub-optimal assignment vs. Optimal assignment.

As Fig. 1(a) shows, in the online setting, task t_1 arrives first and would be assigned to worker w_1 since they are close to each other. Note that such a decision is irrevocable in the online setting. Subsequently, w_1 will no longer be available no matter what future tasks arrive. When task t_2 arrives later, because worker w_1 is occupied, t_2 cannot be assigned. w_1 is the only worker who can perform t_2 , as t_2 is outside the reachability range R_{w_2} of worker w_2 . In contrast, when t_1 and t_2 come in a batch, they can be properly assigned to w_2 and w_1 , respectively.

No existing research has specifically targeted the batch-based privacy-preserving task assignment, however, as a matter of fact, batch-based setting has been widely adopted in the spatial crowdsourcing industry. For example, DiDi [3], the leading car-hailing platform in China, accumulates orders (tasks) per time window, and jointly dispatches them to drivers (workers) [26, 29]. Figure 2 gives an example of our studied problem in this paper. Fig. 2(a) shows the true locations of workers w_1, w_2 and tasks t_1, t_2 . w_1 and w_2 have a reachability range R_{w_1} and R_{w_2} respectively, denoting the maximum distance that they are willing to travel. In Fig. 2(b), the locations are perturbed using Geo-I. Our problem is to assign tasks to suitable workers to maximize the number of assigned tasks subjecting to workers' reachability constraint.




Figure 2: Our Privacy-preserving Batch-based Task Assignment (PBTA) problem.

The obfuscated locations bring technical challenges to the task assignment problem. As shown in Fig. 2(b), if we directly use the perturbed locations as the true locations to perform the task assignment, then task t_2 and t_1 are assigned to worker w_2 and w_1 respectively (denoted by the red arrows). However, according to the true locations, t_2 lies outside the range of w_2 , and t_1 lies outside the range of w_1 . The aforementioned assignment finally has no task successfully executed. In fact, the optimal assignment should be assigning t_2 to w_1 and t_1 to w_2 (denoted by the green arrows in Fig. 2(a)), which would finally execute both tasks.

Table 1: Notation.

Symbol	Description
w, t, W, T	A worker, a task, a set of workers, a set of tasks
w', t'	A worker and a task after perturbation
l_w, l_t	True locations for a worker w and a task t
l_w', l_t'	Perturbed locations for a worker w and a task t
R_w	Reachable distance for a worker w
$d(\cdot)$	Euclidean distance function
M	An assignment (matching) between workers and tasks
M_0	A baseline matching obtained by oblivious baseline methods
p, P	p is a matched worker/task pair $p = (w, t)$. $P = W \times T$ is the set of all possible worker-task pairs.
$l = \epsilon r$	Privacy parameters

To address the aforementioned challenge, we propose an innovative solution k -Switch. It first employs probabilistic inference over the perturbed locations to obtain an initial assignment, then uses an encryption-based secure computation method to improve the assignment quality significantly without compromising privacy. To address the efficiency issue of existing Homomorphic Encryption techniques, it divides workers into small groups of size k , and apply HE within the small groups only. Using the example in Fig. 2, k -Switch uses perturbed locations to obtain a preliminary assignment first, and then allows worker w_1 and w_2 to use secure communication to obtain the *true distance* from the tasks and switch tasks in order to obtain a better assignment.

To summarize, we make the following contributions in this paper.

- We propose k -Switch, an innovative solution to address the challenging Privacy-preserving Batch-based Task Assignment (PBTA) problem. To the best of our knowledge, this is the first work targeting the *batch-based setting* of the privacy-preserving task assignment problem. The PBTA problem is defined in Section 3 together with two proposed oblivious baselines.
- On top of the oblivious baselines, we propose the k -Switch method, which combines probabilistic analysis and encryption-based secure computation techniques to achieve significant assignment quality improvements without compromising privacy. It mitigates the efficiency issue of previous encryption-based secure computation techniques by doing the computation within the small groups. We introduce k -Switch in Section 4.
- We conduct extensive experiments to validate the efficiency and effectiveness of the proposed algorithms. In terms of the number of successfully assigned tasks, the k -switch method assigns up to $5.9\times$ more tasks than other proposed baselines, and assigns up to $1.74\times$ more tasks than the existing online solution. The results are shown in Section 5.

In addition, we cover the preliminaries and necessary background in Sec. 2, compare with other related works in Sec. 6 and conclude the paper in Sec. 7. The notations used in this paper are summarized in Table 1.

2 BACKGROUND

2.1 Task assignments in spatial crowdsourcing

We introduce the task assignment problem, especially the batch-based setting, which is the primary focus of this paper.

Definition 2.1. (Task assignment problem in spatial crowdsourcing [22, 27]) Given a set of workers W and a set of tasks T , the task assignment problem returns an assignment (matching) M of tasks to workers $M = \{(w, t) | w \in W, t \in T\}$ such that for some given objective function $\Psi(\cdot)$, $\Psi(M)$ is optimized (maximized or

minimized):

$$\Psi(M) = \sum_{(w,t) \in M} \psi(w,t),$$

Following the setting in [21], we have the following spatial constraint and objective function. Each worker is willing to travel at most $R_w \in \mathbb{R}$, i.e., for each pair $(w, t) \in M$, $d(l_w, l_t) \leq R_w$. Each successfully assigned task carries a unit utility. The task assignment problem maximizes $\Psi(M) = |M|$, the size of the matching.

In this paper, we focus on the batch-based setting (also referred as offline or static setting): the locations of *all* workers and tasks are known at the beginning. Previous works [18, 21] mainly focus on the privacy preservation for the online setting: tasks arrive one by one, and each task needs to be assigned immediately upon arrival and cannot be re-assigned to other workers no matter what future tasks arrive.

2.2 Geo-indistinguishability

Geo-indistinguishability (Geo-I) [5] extends the traditional and well-adopted privacy notion – differential privacy [8] to location-based systems.

Definition 2.2. (Geo-I) For all true locations x, x' , a privacy parameter ϵ , a mechanism M satisfies ϵ -Geo-I iff:

$$d_\rho(M(x), M(x')) \leq \epsilon d(x, x'),$$

where $d(x, x')$ is the Euclidean distance between x and x' while $d_\rho(M(x), M(x'))$ is the multiplicative distance between two distributions $M(x)$ and $M(x')$. $M(x)$ and $M(x')$ are the distributions of perturbed locations based on the original location x and x' respectively.

One particular mechanism satisfying Geo-I is drawing random noise from the planar Laplace distribution [5]. Given the privacy parameter $\epsilon \in \mathbb{R}^+$, the actual location $x_0 \in \mathbb{R}^2$, the probability density function of a noisy location $x \in \mathbb{R}^2$ is:

$$D_\epsilon(x_0)(x) = \frac{\epsilon^2}{2\pi} e^{-\epsilon d(x_0, x)}, \quad (1)$$

where $\frac{\epsilon^2}{2\pi}$ is the normalization factor.

3 PROBLEM DEFINITION

We formally define the Privacy-preserving Batch-based Task Assignment (PBTA) problem in this section and introduce our proposed oblivious baselines.

3.1 PBTA problem

Definition 3.1. (Privacy-preserving Batch-based Task Assignment (PBTA) problem) Given a set of workers W and a set of tasks T , the perturbed location l_t for each task t , the perturbed location l_w for each task w , the reachable distance R_w for each worker w , the Euclidean distance function $d(\cdot)$, the PBTA problem is to return an assignment (matching) M of tasks to workers $M = \{(w, t) | w \in W, t \in T, d(l_w, l_t) \leq R_w\}$ such that the following objective function is maximized:

$$\Psi(M) = |M|.$$

Note that in the definition above, only the perturbed locations l_w, l_t are available in the input. On the other hand, for the objective

we are maximizing, each pair in the matching M needs to satisfy the spatial constraint w.r.t. the true locations l_w, l_t . We have shown a simple example in Sec. 1 (Figure 2) to illustrate the problem.

In our setting, The perturbed locations of both workers l_w and tasks l_t are obtained by applying Geo-I with the privacy level $l = \epsilon r$. The privacy level is the same for all workers and tasks. Also, one worker w takes at most one task, and each task t only needs to be matched to one worker. The setting follows [21].

3.2 Privacy model

3.2.1 System model. We follow the system model assumptions of [21]. We have three parties in our system: the server, the workers, and the task requesters (short as *tasks* hereafter).

Because the server is untrusted, when workers and tasks submit their locations, they only send the perturbed locations. The server is untrusted, as centralized servers are usually vulnerable to attacks, suffering massive data leak. For example, Facebook security breach exposes 50 million users' data [11].

The task assignment is done at the server. After workers and tasks submit perturbed locations to the server, the server runs its task assignment algorithm, obtains an assignment M , and notifies workers and tasks about their assignments.

3.2.2 Adversary model. Similar to the setting in [21], we adopt a semi-honest model, which assumes that all participating parties (the workers, the task requesters, and the server) are curious but not malicious.

They are curious about the private and sensitive information about other parties. So, we try to prevent sensitive information of any party from being shared with other parties. On the other hand, they are not malicious and follow system protocols. They do not collude with each other to gain extra information.

3.3 Baselines

We propose a simple baseline solution Oblivious-M, which directly uses the perturbed (observed) locations to obtain a matching between workers and tasks. Then, we incorporate probabilistic analysis of reachability between workers and tasks, and propose another solution Oblivious-RR, which is based on randomized rounding.

The term ‘Oblivious’ indicates that both of these methods only access the perturbed locations, and true locations have never been accessed and disclosed in any way when running these methods.

3.3.1 Oblivious-M. The Oblivious-M method is the simplest baseline solution we propose. It builds the reachability graph from the perturbed locations, and runs the Max-Flow algorithm (e.g., Ford-Fulkerson algorithm [7]) to obtain a maximized cardinality matching.

We use a running example in Figure 3 (modified from Fig. 3 of [21]) to demonstrate the basic steps of Oblivious-M. In the PBTA problem, we are given the observed (perturbed) locations of workers and tasks (Fig. 3(a)). In Oblivious-M method, directly using the observed locations, if a task t is within the range R_w of worker w , we create an edge between w and t indicating they are reachable. In this way, we build the reachability graph as shown in Fig. 3(b).

The reachability graph is a bipartite graph with all workers on one side and tasks on the other side, and all edges (indicating reachability) in between. Next, we simply add superficial source/sink

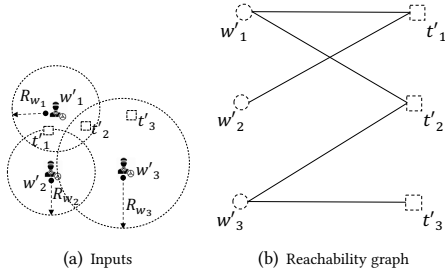


Figure 3: A running example of PETA problem

nodes, edges with capacity 1 between the source node and every worker, and edges with capacity 1 between the sink node and every task to build the flow network (see Fig. 4). It is shown that the maximum cardinality matching from the reachability graph corresponds to the Max-Flow on the constructed flow network. So, we run standard Max-Flow algorithms (such as Ford-Fulkerson [7]) on the constructed flow network, and return all saturated edges (w, t) (with flow value 1 on the edges) as the output M . In our running example, the maximum flow obtained is shown in Fig. 4. The returned matching is $M = \{(w_1, t_2), (w_2, t_1), (w_3, t_3)\}$.




Figure 4: The maximum flow obtained on our running example.

When we measure the size of the matching M , we need to use the true locations to check whether a task is indeed within the reach of the assigned worker. The true locations and the reachability graphs built from the true locations are shown in Figure 5. In our returned matching M based only on the perturbed locations, t_2 and t_1 are indeed within the range of w_1 and w_2 respectively. However, t_3 lies outside the range of w_3 . Thus, the size of the matching we find is 2 (instead of 3).




Figure 5: The true locations of our running example.

For the time complexity, the fattest-path first implementation of Ford-Fulkerson is $O(|E| \cdot opt)$, where E denotes edges for the flow network, and opt is the optimal value of the flow. We let $n = \max(|W|, |T|)$, the larger value of the sizes of the workers and the tasks. In our constructed flow network, assuming all tasks are connected with all workers in the worst case, $|E| = O(n^2)$. As for opt , the optimal flow value is bounded by n . Overall, Oblivious-M has complexity of $O(|E| \cdot opt) \rightarrow O(n^2 \cdot n) \rightarrow O(n^3)$.

3.3.2 *Oblivious-RR*. When we construct the reachability graph in the simplest baseline solution Oblivious-M, an edge between a worker w and a task t is either 1 (if perturbed locations show that




Figure 6: The maximum weight flow and the rounded matching.

t is reachable from w) or 0 (when t is not reachable from w). We make improvements in this step by adopting probabilistic analysis. Intuitively, based on perturbed locations, some worker-task pairs should have higher likelihood that they are indeed reachable, and on the other hand, some other worker-task pairs should have lower likelihood of being reachable.

We continue with the running example in Figure 3 to further explain the intuition of the probabilistic analysis. If we focus on w'_1 , and compare two tasks t'_1 and t'_2 , then based on the perturbed locations, t'_1 looks much closer to w'_1 compared to t'_2 (the distance between w'_1 to t'_1 is almost $1/2$ of t'_2 , as t'_2 is almost at the periphery of the reachable circle). Oblivious-RR incorporates the probabilistic analysis and gives each edge in the reachability graph a fractional weight ranging from 0 to 1, instead of a binary 0/1 value. Based on the fractional flow network, we design randomized techniques to obtain a matching. Due to space limit, we leave the details of Oblivious-RR in the full version [14]. Here, we show a particular matching obtained after the randomized rounding in Fig. 6(b), which assigns the same amount of tasks as the optimal matching.

The time complexity of Oblivious-RR is the same with Oblivious-M, as it only adds a post randomized rounding which takes $O(n^2)$ time, with $n = \max(|W|, |T|)$. Overall, Oblivious-RR runs in $O(n^3)$ time.

4 K-SWITCH

We present the main contribution of this paper – our proposed k -Switch in this section. We introduce the basic idea of the entire solution first, and then present the technical details in each step of the method.

4.1 Overview

k -Switch method is a novel task swapping method, which uses coordination between workers to achieve utility gain. It trades off modest system overhead with secure computation between workers to increase the number of successfully assigned tasks, without privacy leak.

Figure 7 shows the basic idea of the method. k -Switch starts with a baseline matching obtained by our proposed Oblivious baselines, and then makes an improvement on the matching with the following steps:

Step 1) Grouping (Sec. 4.3): carefully groups workers into groups of size k , where k is a small number selected by the server.

Step 2) k -HE (Sec. 4.4): for each k -group, the workers and the task inside the group uses secure computation to calculate the true distance based on encrypted true locations. Then, workers inside




Figure 7: k -Switch

the k -group swap tasks if such swapping achieves utility gain, i.e., increase the number of successfully assigned tasks.

Step 3) λ -Opting (Sec. 4.5): server iterates Step 1) and Step 2), each time grouping workers into small k -groups and letting them use k -HE to communicate inside the group to achieve utility gain.

We introduce the details of each step in the following sections. We show Step 1) Grouping requires solving an NP-hard problem when $k \geq 3$, and our solution includes an efficient greedy grouping algorithm.

4.2 Baseline matching

k -Switch starts with a baseline matching M_0 obtained by either Oblivious-M or Oblivious-RR proposed in Section 3. According to the experiment results, we find that Oblivious-M outperforms Oblivious-RR in terms of the number of successfully assigned tasks, so we use it in k -Switch. In the future, if other better Oblivious matching methods are proposed, they could also be incorporated to find the baseline matching M_0 .

For clarify of presentation, we formally define the baseline matching M_0 .

Definition 4.1. (Baseline matching) For a PBTA problem as defined in Def. 3.1, we use Oblivious-M to find a baseline matching M_0 :

$$M_0 = \{(w, t) | w \in W, t \in T\}$$

Again, to illustrate our idea clearly, we continue with the running example in Figure 4. Using Oblivious-M, we obtain a baseline matching $M_0 = \{(w_1, t_2), (w_2, t_1), (w_3, t_3)\}$. At this stage, only perturbed locations shown in Figure 3 is observable to the server.

As we illustrate previously, this baseline matching could be sub-optimal when we use true locations to verify the reachability constraints of the workers. In the next step Grouping, we divide workers into small groups of size k for further optimization.

4.3 Grouping

4.3.1 Definitions and k -Grouping problem. The purpose of grouping is to group nearby workers and tasks into small groups, such that there are good chances that swapping tasks inside the small group could lead to improvement of the number of successfully assigned tasks (the utility of the matching). The small groups are referred as k -groups. Next, we formally define the k -group and the k -Grouping problem.

For clarify of presentation, we use $P = W \times T$ to denote the set of all possible matched worker/task pairs. For example, in the baseline matching M_0 in our running example, we have three matched pairs,

$p_1 = (w_1, t_2) \in P$, $p_2 = (w_2, t_1) \in P$, $p_3 = (w_3, t_3) \in P$. The baseline matching could also be denoted as $M_0 = \{p = (w, t) | w \in W, t \in T\}$.

Definition 4.2. (k -group) A k -group is a set of k matched worker-task pairs in the baseline matching. A k -group g contains k different workers and their matched tasks:

$$g = \{p = (w, t) | p \in M_0\} \quad s.t. \quad |g| = k.$$

Back to the running example, if $k = 2$, then a k -group becomes a 2-group which contains 2 matched pairs. For example, $g_1 = (p_1, p_2)$ is a possible 2-group, $g_2 = (p_1, p_3)$ is another possible 2-group. For 2-group g_1 , it contains worker from the matched pair p_1 , which is w_1 . It also contains the worker from the matched pair p_2 , which is w_2 . The 2-group g_1 also contains the tasks that are matched with the workers, which are task t_2 in pair p_1 and t_1 in pair p_2 .

The purpose of grouping is to group nearby workers and tasks together such that there are good opportunities that swapping tasks between them leads to better assignment. Some k -groups are better than the others. To evaluate the quality of each k -group, we propose a simple yet effective measurements, verified by experiments, called obfuscation-score (short as OScore). We first define the OScore on a 2-group and then extend it to k -group.

Definition 4.3. (Obfuscation-score) Given a 2-group $g = \{p_1, p_2\}$, we define the obfuscation-score (or OScore) of the 2-group as:

$$OScore(g) = d(l'_{p_1.w}, l'_{p_2.w}) + d(l'_{p_1.t}, l'_{p_2.t}).$$

Here p_1 and p_2 are the two matched pairs inside the 2-group g . Notation $p.w$ and $p.t$ respectively denote the worker and the task from the matched pair.

The OScore definition is straightforward, it measures two distance: i) the distance between the perturbed locations of the two workers in the 2-group; and ii) the distance between the perturbed locations of the two tasks that are assigned to the two workers in the baseline matching M_0 . Then OScore is the sum of the two distance.

Intuitively, if both workers and their assigned tasks appear to be close to each other, they are spatially *clustered*, and should be grouped together in a k -group. Other measures more sophisticated than OScore could be defined and used, however we adopt this OScore as it is very efficient to compute and yet effective, as to be demonstrated in our experiments. Next, we extend the OScore measurement from 2-group to k -group.

Definition 4.4. (OScore of a k -group) For a k -group $g = \{p_1, \dots, p_k\}$, we define its OScore as:

$$OScore(g) = \sum_{1 \leq i < j \leq k} OScore(\{p_i, p_j\}).$$

For a k -group g , its OScore is the sum of the OScore of g 's subsets of size 2. We abuse notations to use the same OScore to refer to different definitions for 2-group and k -group when $k \geq 3$.

Before we introduce the k -grouping problem, we define the term k -division as a collection of k -groups that we select out of the baseline matching M_0 .

Definition 4.5. (k -division) Given a baseline matching $M_0 = \{p | p \in P\}$, a k -division $D = \{g_1, \dots, g_d\}$ is a set of non-overlapping

k -groups of M_0 , where the size of the k -division is $d = \lceil |M_0|/k \rceil$. We define the score of a k -division as:

$$\text{score}(D) = \sum_{1 \leq i \leq d} \text{OScore}(g_i).$$

Remark: i) for $1 \leq i \leq d$, $|g_i| = k$ except for at most one group, when $|M_0|$ is not a multiple of k ; ii) k -groups are non-overlapping (disjoint) when workers from any k -group is different from any other k -groups.

The k -division is a collection of k -groups from the baseline matchings. It divides workers into small k -groups, each with exact size k , except for at most one group. The exception happens when the total number of workers is not a multiple of k , and the exceptional group has a size in the range of $[0, k - 1]$. The score of the k -division is defined as the sum of OScore of all its k -groups.

Next we are ready to formally define the k -Grouping problem. It divides the workers in the baseline matching M_0 into small groups to create a k -division as defined in Def. 4.5, and minimizes the score of the k -division.

Definition 4.6. (k -Grouping problem (KGP)) Given a matching $M_0 = \{p|p \in P\}$, the k -grouping problem returns a k -division D^* with minimal score. Formally, for any other k -division D :

$$\text{score}(D^*) \leq \text{score}(D).$$

Next, we provide theoretical analysis on the hardness of k -Grouping problem (KGP). We show for $k = 2$, KGP is polynomial-time solvable, while for $k \geq 3$, it is hard to approximate.

4.3.2 Algorithms for $k = 2$.

THEOREM 4.7. *When $k = 2$, KGP is in P class, solvable in polynomial time.*

PROOF. The overall idea of the proof is that, when $k = 2$, the KGP is equiv. to finding the maximum weight matching on a general graph. We construct a graph $M_{\mathcal{G}}$ from the baseline matching $M_0 = \{p|p \in P\}$ as follows. For each matched pair $p \in P$, we create a vertex p and add it to graph $M_{\mathcal{G}}$. We create an edge between any two vertices p . It is a complete graph. We define the weight of the edges as follows. For two vertices $p_i \in M_0$ and $p_j \in M_0$, we set the edge weight $w(p_i, p_j) = \text{OScore}(\{p_i, p_j\})$, where $\text{OScore}(\{p_i, p_j\})$ is defined in Def. 4.3.

We could then show that the maximum weight matching on $M_{\mathcal{G}}$ corresponds to the optimal solution to the KGP when $k = 2$. We defer the details of the proof to the full report [14]. \square

THEOREM 4.8. *When $k = 2$, KGP is solvable with time complexity $O(|M_0|^{2.37})$.*

PROOF. For the maximum weight matching problem on a general graph, it is shown that it is among the *hardest* problem that could be solved in polynomial time, with $O(|V|^3)$ time complexity [10], where V denotes the vertex set of the graph. In our setting, $|V| = |M_0|$, the size of the baseline matching.

In our setting, as the constructed graph $M_{\mathcal{G}}$ is a complete graph, and convertible to bipartite graph via a simple *graph transformation* method (please refer to the appendix [14]). There exists matrix multiplication algorithms on the transformed bipartite graph to obtain the maximum weight matching with time complexity $O(|V'|^{2.37})$.

Algorithm 1: Greedy-Grouping

Input: A baseline matching M_0 . Perturbed locations $l_{w'}$ and $l_{t'}$ for each $w \in W$ and $t \in T$.
Output: A k -division D .

- 1 Initialize a heap h , an empty set D
- 2 **foreach** $p_1 \in M_0$ **do**
- 3 **foreach** $p_2 \in M_0$ **do**
- 4 **if** $p_1.w \neq p_2.w$ **then**
- 5 $\text{oscore} := d(l'_{p_1.w}, l'_{p_2.w}) + d(l'_{p_1.t}, l'_{p_2.t})$
- 6 $h.\text{insert}(\{p_1, p_2\}, \text{oscore})$
- 7 $d := \lceil |M_0|/k \rceil$ // Calculate how many k -groups
- 8 **for** $i := 0; i < d; i = i + 1$ **do**
- 9 $g := \{\}$
- 10 **while** $g.\text{size}() < k - 1$ **do**
- 11 $p_1, p_2 := h.\text{pop}()$
- 12 $g.\text{Insert}(p_1, p_2)$
- 13 Mark $p_1.w$ or $p_2.w$ as used
- 14 $D.\text{Insert}(g)$
- 15 **return** D

Because the graph transformation doubles the number of vertices of the original graph $M_{\mathcal{G}}$, we have $|V'| = 2 * |V| = 2|M_0|$. Overall, KGP problem is solvable in $O(|V'|^{2.37}) \rightarrow O(2^{2.37} * |V|^{2.37}) \rightarrow O(|M_0|^{2.37})$. If we let $n = \max(|W|, |T|)$, because $|M_0|$ is bounded by n , then KGP is solvable in $O(n^{2.37})$. \square

4.3.3 Algorithms for $k \geq 3$.

THEOREM 4.9. *When $k \geq 3$, KGP has no polynomial time approximation algorithm with finite approximation ratio unless P=NP.*

PROOF. We show a polynomial reduction of the *Perfectly Balanced Graph Partition* (PBG) problem to the KGP. Because PBGP has no polynomial approximation algorithm with finite approximation ratio unless P=NP [4], our KGP has the same hardness.

We review PBGP: given a graph $G = (V, E)$, with weight $w(e)$ on each edge e . For an integer $p \geq 2$, a p -partition is p disjoint subsets with equal sizes: $V = V_1 \cup \dots \cup V_p$. We assume $|V|$ is a multiple of p here, so $|V_1| = \dots = |V_p| = |V|/p$. The decision version of PBGP is that given a positive integer W , is there a p -partition such that for the cross edges set $E' \subset E$, which have two endpoints in the two different sets V_i , i.e., $E' = \{(v_i, v_j) \in E | v_i \in V_i, v_j \in V_j, i \neq j\}$, the sum of the weights of edges in E' is less or equal to the given integer W , i.e., $\sum_{e \in E'} w(e) \leq W$?

Please refer to the full report [14] for the details of the proof. Because we could show that PBGP \leq_p KGP, and since PBGP has no polynomial time approximation algorithm with finite approximation factor unless P=NP [4], KGP has the same hardness. \square

Because of the intractability of KGP, we propose an efficient greedy method to find a k -division. The greedy algorithm *packs* k -group one by one, each time starting with an empty set, and keeps adding a new 2-group $\{p_i, p_j\}$ with smallest OScore to the current k -group. If k is an odd number, it randomly picks the last matched pair p . The detailed steps is shown in Alg. 1.

At Line 6, the algorithm uses a heap storing all combinations of pairs in the baseline matching $\{p_i, p_j\}$ with its OScore as the sorted




Figure 8: k -HE protocol (the case for $k = 4$)

key. Every time we could pop the pair $\{p_i, p_j\}$ with the smallest OScore and add them to the current k -group. After each k -group is formed, we continue to the next one until a k -division is obtained and returned. Line 10-13 greedily add two matched pairs to the current k -group g , until its size reaches k . In total, we form d groups, as calculated at Line 7. We omit some details for checking used workers and randomly picking the last item (see full version [14]).

The time complexity of Alg. 1 is dominated by the heap construction, which takes $O(e \log e)$, where e is the total number of elements inserted to the heap. We know e is all the combinations of matched pairs, as shown at Line 2-3, so $e = O(n^2)$, where $n = \max(|W|, |T|)$, the size of the baseline matching. In conclusion, for Alg. 1 has an $O(n^2 \log n^2) \rightarrow O(n^2 \log n)$ time complexity.

4.4 k -HE protocol

The purpose of the previous step Grouping is to divide all workers into small groups of size k . In this section, we introduce how workers inside each small k -group utilize secure computation in parallel to increase the number of successfully assigned tasks via task swapping.

Our k -HE protocol runs in the small group of size k . The secure computation is based on the Paillier Crypto-system [16] in Homomorphic Encryption (HE), which allows Homomorphic Addition and Homomorphic Multiplication. Similar to the global task assignment setting in HESI framework [15], we also use HE for secure distance calculation. In contrast, our protocol is restricted to small size k (ranging from 2 to 8). We allow workers and tasks inside the small group to communicate the encrypted true locations with one another, and if the number of successfully assigned tasks could be improved based on their true locations, then workers swap tasks between themselves.

Figure 8 gives an illustration of the protocol. Two entities (either worker or task) out of the k -group are randomly elected and serve as the proxy servers P_a and P_b . Then, we perform secure distance calculation between each pair of workers and tasks following the major steps of HESI [15] (details deferred to our full report [14]). As for the time complexity, two proxy servers enumerate all combinations of worker-task pairs in the k -group and compute the true distance, with $O(\binom{k}{2}) \rightarrow O(k^2)$ time. Then, P_b runs a matching algorithm w.r.t. the true distances, with $O(k^3)$ time using the max-flow algorithm similar to the Oblivious-M baseline (Sec. 3.3.1). The overall time complexity is $O(\binom{k}{2} + k^3) \rightarrow O(k^3)$.

4.5 λ -Opting

The previous section describes how small k -groups execute k -HE protocol in parallel and workers swap tasks if task swapping increases the number of successfully assigned tasks within the group.

The last phase of k -Switch is λ -Opting, which iterates Grouping and k -HE protocol for λ rounds. It also stops if no utility gain is obtained at the current round. λ is a system parameter controlling the trade-off between utility gain and system overhead.




Figure 9: λ -Opting.

Each round of λ -Opting executes the grouping and the k -HE protocol, taking $O(k^3 + n^2 \log n)$ time. The time complexity of k -Switch is thus $O(\lambda k^3 + \lambda n^2 \log n)$, where $n = \max(|W|, |T|)$.

5 EXPERIMENTAL STUDY

We conduct extensive experiments on both the real-world dataset and the synthetic dataset to validate the effectiveness and efficiency of our proposed k -Switch method.

With respect to the effectiveness, as measured by the number of successfully assigned tasks, experiments show that k -Switch assigns up to 5.9 \times more tasks than other batch-based baselines, and assigns up to 1.74 \times more tasks than the competing online method SCGuard. In terms of efficiency, as measured by the running time, our method is efficient, finishing within 1.5 minutes on datasets of moderate sizes (500 workers and 500 tasks). While being slightly slower than other methods (slower than SCGuard by about a constant factor of 2), k -Switch is considered cost-effective because it trades off minor system overhead with considerable utility gain.

5.1 Experimental setup

5.1.1 Datasets. We conduct the experiments on both the real-world and the synthetic dataset. The real-world dataset is the taxi dataset from Didi Chuxing [26]. For the synthetic dataset, we randomly sample workers and tasks' locations from the range $[0, 8000] \times [0, 8000]$.

5.1.2 Baselines. The baselines we test include the two baseline solutions we propose: **Oblivious-M** (short as OM, introduced in Sec. 3.3.1) and **Oblivious-RR** (ORR, Sec. 3.3.2). In addition, we test **SCGuard** (SCG, [21]), an online method allowing each newly arrived task to interactively check several other workers to see whether the task could be assigned.

5.1.3 Metrics and control variables. **Control variables:** Number of workers $w \in [100, 200, 500, 1000]$. Number of tasks $t \in [100, 200, 500, 1000]$. System parameter $k \in [2, 4, 6, 8]$, $\lambda \in [5, 10, 20]$. Privacy requirement $\epsilon \in [0.4, 1.25, 2.5]$. We set $r = 1000m$ as constant. Varying ϵ corresponds to varying Geo-I privacy level $l = \epsilon r \in [400, 1250, 2500]$. Default parameters are in boldface.

Metrics. We focus on 1) effectiveness (*utility*), measured by the number of tasks assigned, and 2) efficiency, measured by running time in seconds.

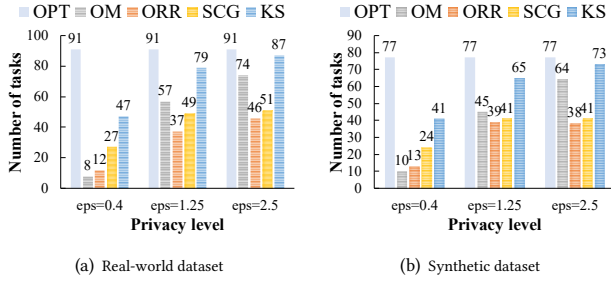


Figure 10: Number of tasks assigned for different methods, over different privacy levels, 100 workers vs. 100 tasks, $k = 2$, $\lambda = 20$.

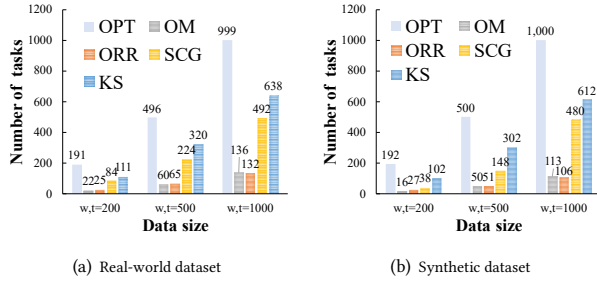


Figure 11: Number of tasks assigned, on different input sizes (200-1000), $\epsilon = 0.4$, $k = 2$, $\lambda = 20$.

System configuration. The experiments are performed on a MacBook Pro with 1.4GHz Quad-Core Intel Core i5 and 16GB 2133MHz LPDDR3 memory, running MacOS 11.0. The methods were implemented in Python.

5.2 Experimental results

5.2.1 Effectiveness. Overview: k -Switch (short as KS in the figures) outperforms other baselines by significant margins, over different privacy levels (Fig. 10), across datasets of different sizes (Fig. 11). For a stricter privacy parameter $\epsilon = 0.4$ on the taxi dataset (shown in Fig. 10(a)), k -Switch achieves $5.9\times$ improvement over baseline Oblivious-M, and $1.74\times$ improvement over the competing online method SCGuard.

Details of results. Fig. 10 shows the number of successfully assigned tasks obtained by different methods, over different privacy levels on a dataset of 100 workers vs. 100 tasks. We compare the optimal matching with the Oblivious baseline methods. The optimal matching (OPT) is obtained by using the ground-truth locations, which are not available in the inputs to our PBTA problem. The Oblivious baselines OM and ORR are using only the perturbed locations. At privacy level $\epsilon = 0.4$ (shown in Fig. 10(a)), the gap between OPT and OM is 82. The OPT is about 11 times larger than OM. This validates the motivation of our research: privacy-preserving techniques perturb the locations of workers and tasks, and directly assigning tasks based on perturbed locations is erroneous and prone to sub-optimal assignment. The gap is also observable on the synthetic dataset, shown in Fig. 10(b).

Varying privacy levels: k -Switch outperforms all other methods on different privacy levels, assigning 47 tasks for $\epsilon = 0.4$ on taxi data (shown in Fig. 10(a)), achieving $5.9\times$ improvement over OM, which assigns 8 tasks. It also achieves $1.74\times$ improvement over the SCGuard, which assigns 27 tasks. When ϵ gets larger, the privacy requirement gets less strict, the gap between the OPT and the oblivious OM and ORR gets smaller. For $\epsilon = 2.5$, OM assigns 74 tasks,

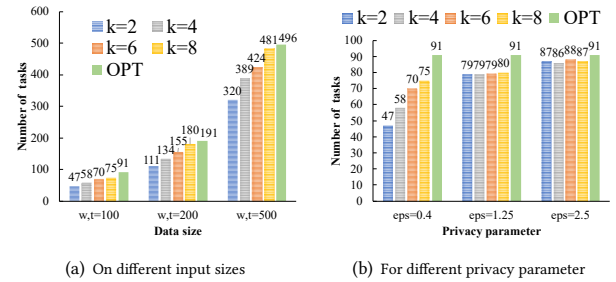


Figure 12: Effect of parameter k for k -Switch method on the number of tasks assigned.

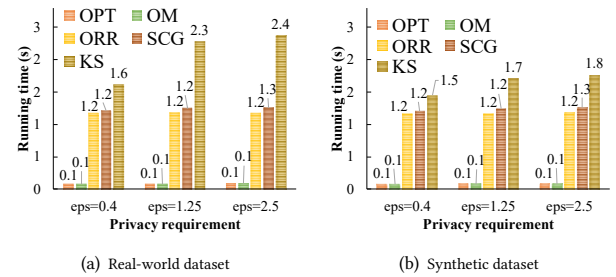


Figure 13: Running time (seconds) for different methods, over different privacy levels, 100 workers vs. 100 tasks, $k = 2$.

much closer to the 91 tasks from the OPT solution, as compared to a stricter ϵ . Nevertheless, k -Switch delivers strong performance, and assigns 87 tasks, which is close to the OPT.

Varying data size: Fig. 11 shows the number of successfully assigned tasks of different methods, on datasets of different sizes. First, as the number of worker and task increases from 200 to 1000, the OPT result increases from 191 to 999 on taxi dataset (Fig. 11(a)), and 192 to 1000 on the synthetic dataset (Fig. 11(b)). The gap between OPT and OM is consistently large. For the taxi dataset (Fig. 11(a)), the OPT/OM ratio is $191/22 = 8.7$ for $w, t = 200$ and $999/136 = 7.3$ for $w, t = 1000$, respectively. k -Switch significantly improves over the baseline OM and the online SCGuard method. It achieves $5.05\times$ improvement over OM for $w, t = 200$, $5.33\times$ for $w, t = 500$, and $5.64\times$ for $w, t = 1000$. As for the comparison with SCGuard, k -Switch obtains $1.32\times$, $1.43\times$, and $1.30\times$ improvement for $w, t = 200, 500$, and 1000 respectively. The same behavior is observed on the synthetic dataset (Fig. 11(b)).

Varying parameter k : Fig. 12 shows the number of successfully assigned tasks, over different k , on datasets of different sizes (Fig. 12(a)) and over different privacy requirements ((Fig. 12(b))). As Fig. 12(a) shows, when k increases, the number of successfully assigned tasks increases. For the smallest data size, the utility increases from 47 to 75 tasks, as k increases from 2 to 8. On the other hand, on the data size $w, t = 100$, when we vary privacy parameters, the effect of k is not as significant (Fig. 12(b)).

Varying parameter λ : Due to space limit, we defer the results about varying the system parameter λ to the full version [14]. The results verify that λ achieves a tradeoff between the utility and efficiency for k -Switch, which is consistent with our system design.

5.2.2 Efficiency. Overview. While we expect k -Switch to be slower than other methods as the design trades off moderate system overhead with significant utility gain, experimental results show it is only slightly slower than other methods. On our default setting

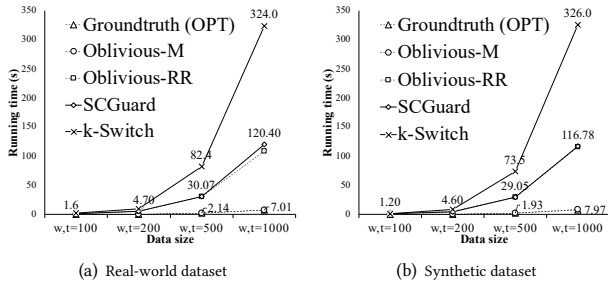


Figure 14: Running time (seconds) on different input sizes.

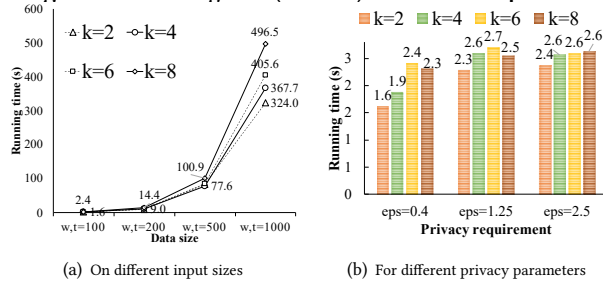


Figure 15: Running time (seconds) of k -Switch for different k .

(Fig. 13), 100 workers vs. 100 tasks, it takes only around 2 seconds to run on a laptop machine. When tested on larger sizes dataset (Fig. 14), k -Switch shows stable efficient running time, a small constant factor slower than the competing SCGuard. Experiments verify our time complexity analysis of the methods.

Details of results. Fig. 13 shows the running time of different methods (in seconds), over different privacy parameters, on datasets of 100 workers vs. 100 tasks. k -Switch obtains around 2 seconds running time on the taxi dataset (Fig. 14(a)), and strictly less than 2 seconds on the synthetic dataset (Fig. 14(b)). The running time is consistent across all privacy levels. When compared with the competing online method SCGuard, k -Switch is slightly slower, with a constant factor of 2.

Varying data size: We look at the asymptotic growth of the running time across different sizes of datasets in Fig. 14. First, the results verify the quadratic time complexity ($O(\lambda k^3 + \lambda n^2 \log n)$, see Sec. 4.5) in terms of n , where $n = \max(|W|, |T|)$, the input size. When the input size doubles, from $w = 500$ to $w = 1000$, the running time increases about 4 times, from 82.4 seconds to 324 seconds on the taxi dataset (Fig. 14(a)). When compared with SCGuard, SCGuard has the same quadratic time complexity, so it also increases about 4 times, from 30.02 to 120.4 seconds when $w = 500$ increases to $w = 1000$. k -Switch is about a 2 times constant factor slower than SCGuard. The same trend is observed on the synthetic dataset (Fig. 14(b)).

Varying parameter k and ϵ : Fig. 15 shows the running time of k -Switch, over different k , on datasets of different sizes (Fig. 15(a)) and for different privacy parameters (Fig. 15(b)). As expected, when the data size is fixed, the running time of k -Switch is larger when k is increased, while the increase is not significant. When we vary the privacy parameter ϵ , the running time of k -Switch is stable, across different values of k (Fig. 15(b)).

6 RELATED WORK

For privacy-preserving task assignment in spatial crowdsourcing, we have discussed the most directly related online methods [18, 21]

in Sec. 1. Here, we include more related works in a broader context of privacy-preserving spatial crowdsourcing.

Different protection methods. Encryption-based techniques have been used to compute the exact assignment between workers and tasks [15]. The computational cost of such pure encryption-based techniques is high and prohibitive for real-world applications. Other privacy protection technique, such as cloaking, is used to protect locations of workers [17], but cloaking is considered as a weaker privacy-preserving technique than Geo-I (Sec. 2.2), because its assumption on adversaries' prior knowledge. There are other related works using Geo-I as the privacy standard [19, 20]. However we adopt a stricter privacy model that the server is untrusted from all workers and tasks, and locations of both parties need to be perturbed before released to any other parties. Similar to our work, the batch-based (offline) setting has also been considered in [23]. However only the workers are protected.

Other crowdsourcing setting. There are other related works in the spatial crowdsourcing spectrum [24, 25, 28]. Different from the task assignment problem, data publishing has been considered in [25]. The truthful rather than privacy-preserving task assignment is considered in [28]. Privacy-preserving crowd-sensing is considered in [24], and the focus is to protect the locations of workers when they report their sensing results, rather than considering our task assignment setting, where workers need to move to a specified location of the assigned task, and both locations (tasks and workers) are perturbed with differential privacy.

7 CONCLUSION

In this work, we target the Privacy-preserving Batch-based Task Assignment (PBTA) problem, where both workers and tasks use Geo-I to perturb their locations before sending them to the untrusted the server. We propose a novel solution k -Switch, which divides the workers into small groups, and uses a secure computation protocol k -HE for inner group communication. If workers inside the groups find that switching tasks between them improves the number of successfully assigned tasks, they swap tasks. Extensive experiments demonstrate that k -Switch is both effective and efficient, achieving significant utility gains with reasonable system overhead.

ACKNOWLEDGMENTS

Libin Zheng's work is supported by the National Natural Science Foundation of China No. 62102463 and the Basic and Applied basic Research Project of Guangzhou basic Research Program 202102080401. Peng Cheng's work is sponsored by the National Natural Science Foundation of China No. 62102149 and Shanghai Pujiang Program 19PJ1403300. Maocheng, Jiachuan and Lei Chen's work is partially supported by National Key Research and Development Program of China Grant No. 2018AAA0101100, the Hong Kong RGC GRF Project 16209519, RIF Project R6020-19, AOE Project AoE/E-603/18, Theme-based project TRS T41-603/20R, China NSFC No. 61729201, Guangdong Basic and Applied Basic Research Foundation 2019B151530001, Hong Kong ITC ITF grants ITS/044/18FX and ITS/470/18FX, Microsoft Research Asia Collaborative Research Grant, HKUST-NAVER/LINE AI Lab, Didi-HKUST joint research lab, and HKUST-Webank joint research lab grants. Xuemin Lin's work is supported by ARC DP200101338. Corresponding author: Peng Cheng.

REFERENCES

- [1] 2021. <https://foursquare.com>.
- [2] 2021. <https://gigwalk.com>.
- [3] 2021. <https://www.didiglobal.com/about-didi/about-us>.
- [4] Konstantin Andreev and Harald Räcke. 2006. Balanced Graph Partitioning. *Theory Comput. Syst.* 39, 6 (2006), 929–939.
- [5] Miguel E. Andres, Nicolas E. Bordenabe, Konstantinos Cjhatzikokolakis, and Catuscia Palamidessi. 2013. Geo-indistinguishability: differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer and communications security*, 901–914.
- [6] Zhao Chen, Rui Fu, Ziyuan Zhao, Zheng Liu, Leihao Xia, Lei Chen, Peng Cheng, Caleb Chen Cao, and Yongxin Tong. 2014. gMission: A General Spatial Crowdsourcing Platform. *Proceedings of the VLDB Endowment* 7, 13 (2014).
- [7] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2001. *Introduction to Algorithms* (2nd ed.). The MIT Press.
- [8] Cynthia Dwork. 2006. Differential Privacy. In *33rd ICALP (Lecture Notes in Computer Science, Vol. 4052)*, Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener (Eds.). Springer, 1–12. https://doi.org/10.1007/11787006_1
- [9] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3-4 (2014), 211–407.
- [10] Jack Edmonds. 1965. Paths, trees, and flowers. *Canadian Journal of Mathematics* 17, 449–467 (1965).
- [11] Mike Isaac and Sheera Frenkel. 2018. Facebook Security Breach Exposes Accounts of 50 Million Users. <https://www.nytimes.com/2018/09/28/technology/facebook-hack-data-breach.html>.
- [12] Leyla Kazemi and Cyrus Shahabi. 2012. GeoCrowd: enabling query answering with spatial crowdsourcing. In *Proceedings of the 20th International Conference on Advances in Geographic Information Systems (SIGSPATIAL)*, 189–198.
- [13] Seon Ho Kim, Ying Lu, Giorgos Constantinou, Cyrus Shahabi, Guanfeng Wang, and Roger Zimmermann. 2014. Mediaq: mobile multimedia management system. In *Proceedings of the 5th ACM Multimedia Systems Conference*. ACM, 224–235.
- [14] Maocheng Li, Jiachuan Wang, Libin Zheng, Han Wu, Peng Cheng, Lei Chen, and Xuemin Lin. 2021. Privacy-Preserving Batch-based Task Assignment in Spatial Crowdsourcing with Untrusted Server [Technical report]. *CoRR* abs/2108.09019 (2021).
- [15] Bozhong Liu, Ling Chen, Xingquan Zhu, Ying Zhang, Chengqi Zhang, and Weidong Qiu. 2017. Protecting Location Privacy in Spatial Crowdsourcing using Encrypted Data. In *EDBT*. OpenProceedings.org, 478–481.
- [16] Pascal Paillier. 1999. Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. In *EUROCRYPT (Lecture Notes in Computer Science, Vol. 1592)*. Springer, 223–238.
- [17] Layla Pournajaf, Li Xiong, Vaidy S. Sunderam, and Slawomir Goryczka. 2014. Spatial Task Assignment for Crowd Sensing with Cloaked Locations. In *MDM (1)*. IEEE Computer Society, 73–82.
- [18] Qian Tao, Yongxin Tong, Zimu Zhou, Yexuan Shi, Lei Chen, and Ke Xu. 2020. Differentially Private Online Task Assignment in Spatial Crowdsourcing: A Tree-based Approach. In *ICDE*. IEEE, 517–528.
- [19] Hien To, Gabriel Ghinita, Liyue Fan, and Cyrus Shahabi. 2017. Differentially Private Location Protection for Worker Datasets in Spatial Crowdsourcing. *IEEE Trans. Mob. Comput.* 16, 4 (2017), 934–949.
- [20] Hien To, Gabriel Ghinita, and Cyrus Shahabi. 2014. A Framework for Protecting Worker Location Privacy in Spatial Crowdsourcing. *Proc. VLDB Endow.* 7, 10 (2014), 919–930.
- [21] Hien To, Cyrus Shahabi, and Li Xiong. 2018. Privacy-Preserving Online Task Assignment in Spatial Crowdsourcing with Untrusted Server. In *Proceedings of the 34th IEEE International Conference on Data Engineering (ICDE)*, 833–844.
- [22] Yongxin Tong, Zimu Zhou, Yuxiang Zeng, Lei Chen, and Cyrus Shahabi. 2019. Spatial crowdsourcing: a survey. *The VLDB Journal* (2019). <https://doi.org/10.1007/s00778-019-00568-7>
- [23] Leye Wang, Dingqi Yang, Xiao Han, Tianben Wang, Daqing Zhang, and Xiaojuan Ma. 2017. Location Privacy-Preserving Task Allocation for Mobile Crowdsensing with Differential Geo-Obfuscation. In *WWW*. ACM, 627–636.
- [24] Leye Wang, Daqing Zhang, Dingqi Yang, Brian Y. Lim, Xiao Han, and Xiaojuan Ma. 2020. Sparse Mobile Crowdsensing With Differential and Distortion Location Privacy. *IEEE Trans. Inf. Forensics Secur.* 15 (2020), 2735–2749.
- [25] Zhibo Wang, Xiaoyi Pang, Yahong Chen, Huajie Shao, Qian Wang, Libing Wu, Honglong Chen, and Hairong Qi. 2019. Privacy-Preserving Crowd-Sourced Statistical Data Publishing with An Untrusted Server. *IEEE Trans. Mob. Comput.* 18, 6 (2019), 1356–1367.
- [26] Zhe Xu, Zhixin Li, Qingwen Guan, Dingshui Zhang, Qiang Li, Junxiao Nan, Chunyang Liu, Wei Bian, and Jieping Ye. 2018. Large-Scale Order Dispatch in On-Demand Ride-Hailing Platforms: A Learning and Planning Approach. In *KDD '18 Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 905–913.
- [27] Yuxiang Zeng, Yongxin Tong, Lei Chen, and Zimu Zhou. 2018. Latency-Oriented Task Completion via Spatial Crowdsourcing. In *ICDE*, 317–328.
- [28] Dongjun Zhai, Yue Sun, An Liu, Zhixu Li, Guanfeng Liu, Lei Zhao, and Kai Zheng. 2019. Towards secure and truthful task assignment in spatial crowdsourcing. *World Wide Web* 22, 5 (2019), 2017–2040.
- [29] Libin Zheng, Lei Chen, and Jieping Ye. 2018. Order Dispatch in Price-aware Ridesharing. *Proc. VLDB Endow.* 11, 8 (2018), 853–865.