3-2017

Privacy in context-aware mobile crowdsourcing systems

Thivya KANDAPPU
Singapore Management University, thivyak@smu.edu.sg

Archan MISRA
Singapore Management University, archanm@smu.edu.sg

Shih-Fen CHENG
Singapore Management University, sfcheng@smu.edu.sg

Hoong Chuin LAU
Singapore Management University, hclau@smu.edu.sg

Citation

KANDAPPU, Thivya; MISRA, Archan; CHENG, Shih-Fen; and LAU, Hoong Chuin. Privacy in context-aware mobile crowdsourcing systems. (2017). *2017 IEEE International Conference on Pervasive Computing and Communications Workshops, PerCom Workshops: Kona, Big Island, Hawaii, March 13-17*. 231-236.
Available at: https://ink.library.smu.edu.sg/sis_research/3630

Privacy in Context-aware Mobile Crowdsourcing Systems

Thivya Kandappu, Archan Misra, Shih-Fen Cheng and Hoong-Chuin Lau
School of Information Systems, Singapore Management University
{thivyak, archanm, sfcheng, hclau}@smu.edu.sg

***Abstract*—Mobile crowd-sourcing can become as a strategy to perform time-sensitive urban tasks (such as municipal monitoring and last mile logistics) by effectively coordinating smartphone users. The success of the mobile crowd-sourcing platform depends mainly on its effectiveness in engaging crowd-workers, and recent studies have shown that compared to the pull-based approach, which relies on crowd-workers to browse and commit to tasks they would want to perform, the push-based approach can take into consideration of worker's daily routine, and generate highly effective recommendations. As a result, workers waste less time on detours, plan more in advance, and require much less planning effort. However, the push-based systems are not without drawbacks. The major concern is the potential privacy invasion that could result from the disclosure of individual's mobility traces to the crowd-sourcing platform. In this paper, we first demonstrate specific threats of continuous sharing of users locations in such push-based crowd-sourcing platforms. We then propose a simple yet effective location perturbation technique that obfuscates certain user locations to achieve privacy guarantees while not affecting the quality of the recommendations the system generates. We use the mobility traces data we obtained from our urban campus to show the trade-offs between privacy guarantees and the quality of the recommendations associated with the proposed solution. We show that obfuscating even 75% of the individual trajectories will affect the user to make another extra 1.8 minutes of detour while gaining 62.5% more uncertainty of his location traces.**

I. Introduction

Mobile crowd-sourcing is an emerging paradigm where tasks requiring the physical presence of a worker at specific locations are distributed to, and completed by, a community of voluntary individuals. Examples of these tasks include, but not limited to, citizen sensing [1], [2], campus monitoring [3], [4], environment monitoring [5], [6], and crowd logistics (e.g., Amazon Flex[1] and Courex Singapore[2]). Such mobile crowd-sourcing platform's success depends mainly on its effectiveness in engaging crowd-workers, which can be done via either the "pull" or "push" approaches. The pull approach, which is widely adopted in almost all commercial systems, relies on crowd-workers to browse and commit to tasks they would want to perform (usually with distance-based filtering or sorting). The push approach, on the other hand, proactively recommends tasks to a worker that the platform believes will best cater to her daily routines. Prior work has demonstrated [3], [4] that, when provided with adequate individual mobility traces, push-based systems can significantly increase the productivity of workers: they waste less time on detours and are able to plan more in advance .

However, the effectiveness of the push-based mobile crowd-sourcing platform depends mainly on the knowledge of individual's location traces, and releasing them continuously to a third-party raises privacy concerns [7], [8], [9] to the users. The continuous release of a user's whereabouts can reveal potentially sensitive details about a user's profile (e.g., her health or relationship status), and even used for stalking. Hence, supporting user location privacy is a desirable goal for push-based mobile crowd-sourcing, as else workers may be reluctant to embrace the platform.

Prior solutions for privacy, provided in the context of location based services [8], [9] find nearest points of interest (such as restaurants and supermarkets) without revealing the actual user locations. However, for mobile crowd-sourcing, detour-minimizing recommendations obviously depend on having accurate knowledge of an individual user's trajectory. In addition, the threat may come from the crowd-sourcing platform itself, which can utilize the user location profiling for its own benefits. Prior solutions that have recommended the use of a trusted proxy fail on this account. Such solutions include k-anonymity [8], [10], data generalization [11], use of cryptography tools [12] and differential privacy [13], [14], [15]. Therefore, we desire a method that does not anonymize users, but rather enables "adequate" trajectory obfuscation, without having to trust any external platforms. We are thus particularly interested on the impact and mechanisms for location privacy under the push-based trajectory aware mobile crowd-sourcing model.

In this paper we propose a mechanism for protecting privacy of workers' locations, where the users obfuscate sensitive locations independently, at-source. The crowd-sourcing server then has access only to the noisy, obfuscated data. Since the server has to generate task recommendations for workers (by matching their daily routine path), data perturbation should try to preserve the platform' salient objectives: minimizing the travel overheads experienced by the workers, and effectively assigning tasks to achieve desired task completion rate . The privacy and performance goals are thus conflicting ones, as excessive data perturbation may convert a user's actual trajectory to a completely distant one, resulting in inappropriate task recommendations.

Our main contributions are:

[1] https://flex.amazon.com/

[2] https://www.courex.com.sg/

- We first demonstrate the specific threats of sharing user locations in the context of push-based mobile crowd-sourcing. We use raw mobility traces of students belong to our urban Singapore Management University (SMU) campus to infer some personal details: matriculation year and school of study. After demonstrating that we achieve high precision in inferring these personal details (83% and 87%, respectively) we propose a mechanism that doesn't require the users to trust any third party while obfuscating the sensitive locations at the user-end.
- We conduct an extensive set of experiments on mobility traces at SMU, and show that the simple location obfuscation technique proposed is able to provide location privacy guarantees without significantly degrading the utility of the platform.

II. Related Work

Mobile crowd-sourcing has recently become a popular approach for executing a variety of location-specific tasks in urban environments. The commercial operators, such as FieldAgent (www.fieldagent.net), GigWalk (www.gigwalk.com), and NeighborFavor (www.favordelivery.com) have already been using mobile crowd-sourcing, where, users *pull* tasks by browsing through the entire pool sorted by proximity to their current locations. *Push* based models [16], [17], in contrast, utilize history-based predictions of a user's likely movement behavior to recommend tasks that lie along, or close to, such movement paths.

Driven by concerns on the consequences of privacy breaches, past work has looked at mechanisms to share location data while preserving privacy. The early works in this area used cryptographic tools [18] to retrieve the information in a private way only between the authorised parties,while barring the access to any other third parties. Masking the locations by using spatial *k*-anonymity [8], [9] (where the location of a user is hidden among (*k-1*) other users who are sharing similar trajectories) was also proposed. The above mentioned techniques use a trusted third party in the middle, to encrypt/anonymize the user locations. A major threat with these approaches is that the trusted third party in the middle becomes a single point of attack. Alternative methods such as privately assigning spatial tasks [19] and private data collection [20] are also proposed; however they fail to consider the impact of such techniques on the accuracy of the recommendations being made.

In recent years, the concept of differential privacy has been explored for location-based applications [14], [15], [13]. The key idea here is to mask a user's actual location to an obfuscated location by applying Laplace noise chosen based on predefined parameters. Very recently, an interactive toolbox called PrivGeoCrowd [21] was proposed to make the "choosing the appropriate parameters to achieve differential privacy" process more intuitive. In our work, we consider push-based crowd-sourcing mechanism, where continual update of user locations is essential to effectively recommend tasks that lie in close proximity to the user's daily movement trajectories. Our main focus is on providing enough privacy protection guarantees to the users while generating meaningful task assignments. In this paper, we adopted a simple perturbation technique that has far better assignment precision even with the perturbed trajectories.

III. System Model

A. Crowdsourcing System

We have evaluated our privacy preserving technique in the context of *TA$Ker*, a mobile crowd-sourcing application that we have developed and deployed in Singapore Management University (SMU). *TA$Ker* is designed for both Android and iOS users and enables them to report the status of various facilities management related resources (e.g., restroom cleanliness, stock level of vending machines, just to name a few).

The first version of *TA$Ker* (with 80 participants over a trial period of 4 weeks) was previously described in [3]. It demonstrated that a push-based platform, which is centrally-coordinated and trajectory-aware, can produce much better results than the conventional pull-based model. More specifically, we showed how its trajectory-aware task recommendation engine can accurately predict worker trajectories (even though the location tracking system is not fine-grained) via the estimation of key *reference locations* (the locations where users stayed for longer). We demonstrated that such a recommendation strategy is effective, even when task execution time windows were limited to 30 minutes: guided users incur only 3 minutes of detour on average, compared to 6 minutes of detour from the alternative strategies. However, as predicted trajectories are used to recommend tasks, the platform did not provide location privacy to users.

B. Trajectories

We utilize user's historical mobility traces to predict their trajectories for three distinct 3-hour time windows daily: (a) 9am-12pm, (b) 12pm-3pm, and (c) 3pm-6pm. Mobility traces are presented as a series of tuples in the form of $< parID, locID, timestamp >$, where $parID$ is an anonymized device ID and $locID$ is a location coordinate represented in the form of $< Buiding, Floorlevel, LandmarkID >$ (with 3 meters of location granularity). The location coordinate is obtained using a Wi-Fi fingerprinting technique, similar to the approach [22] (medium-grained granularity with a median accuracy of 6-8 meters and update frequency of 2-4 minutes).

For each 30-minute time segment, we extract *reference locations* – location where a user spends most time at. For each 3-hour time window, a route is constructed by connecting these reference locations. To account for the stochastic nature of user's movement patterns, we maintain k most probable routes for each 3-hour time window. The actual walking path between consecutive reference locations is obtained by applying a standard shortest path finding algorithm. These predicted walk paths are subsequently used to generate task recommendations to individual workers to minimize expected detours/travel overheads.

Our goal is to prevent any honest-but-curious adversary from extracting sensitive information in a context-aware (push-based) crowd-sourcing system, where user trajectories are used to generate efficient task recommendations. In this paper, we propose to randomly mask a user's *sensitive* location reports (to be defined later) to ensure location privacy. To achieve this, we have to provide answers to the following questions:

- Given trajectory traces from a user, how can an attacker identify sensitive information about particular workers?
- How do we quantify the privacy loss/gain of the user trajectories after they have been masked?

C. *Performance Metrics*

We focus on the following performance metrics:

- **Average Detour:** Since the locations are obfuscated, it will be a challenge for the platform to accurately predict the detour that the user has to make (i.e., distance between the task and user). This will likely increase the deviation that the worker needs to make from her true movement path, to perform the recommended tasks. Our goal here is to ensure that user is not exceeding his travel budget to execute the tasks that lie farther from his routine path.
- **Worker Productivity:** Due to the location data uncertainty, the server may recommend tasks that lie too far from his real routine path. This can cause a worker to perform fewer tasks before her travel overhead budget is depleted, resulting in lower rewards earned.

IV. Data Perturbation for Trajectories

In this section we describe our approach. First we show how an adversary can use the mobility traces data to infer sensitive details about them. Next, we present how our approach minimizes the risk of being physically identified by an attacker. Finally, we present how the effectiveness and the efficiency of the mobile crowd-sourcing system is affected by employing privacy preserving techniques.

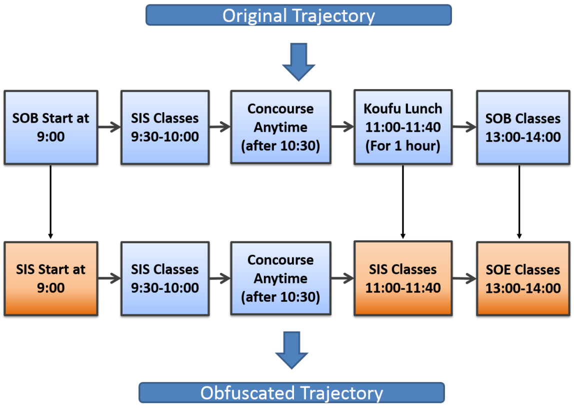

Fig. 1. Perturbing Trajectories

A. *Exposing Privacy Loss*

To illustrate how easily user location privacy is compromised in mobile crowd-sourcing platforms, we consider an honest but curious adversary that follows the system protocol, but is curious to analyse the user trajectories. The adversary can be a server (compromised) that maintains the raw location traces and/or trajectories in its data storage. The adversary is curious to extract additional information from the data to estimate workers' frequent trajectories and sensitive locations.

We illustrate such potential privacy risks by considering the mobility traces (landmark level granularity) and the predicted trajectories of students of our SMU campus. By considering spatio-temporal stay patterns of the students, we show that an adversary can infer an individual's school of study and matriculation year with high accuracy.

Over the period of observation, for each student, we obtain their stay locations (i.e., in locations where a student stayed more than 5 minutes) as a tuple (u,l,t), where u is the anonymised ID of the student, l is the location ID and t is the time stamp. In addition, we leverage the publicly available SMU class timetable. Each individual entry in this dataset is associated with a tuple <*course ID, location, class day, start time, end time*> – this corpus of classes are matched against student's spatio-temporal stay episodes, and we are able to predict the school of study and matriculation year with 87% and 83% precision and 74% and 73% recall. This is an illustration of the type of profile information that may be inferred, simply from a cursory analysis of the physical movement traces of crowd-workers.

B. *Entropy as a Measure of Privacy*

Obfuscation Technique: We use a simple and natural technique by which the user trajectories are obfuscated before storing in the server. For each most sensitive location that is part of user trajectories, a *cloaked location* is substituted that lies within certain radius r from the original location. The radius r will be adjusted to provide different levels of privacy.

Example 1. Consider a trajectory T_i of user U_i, with locations <*library, **classroom X of build. A**, food court, classroom Y of build. B*> A reasonable selection of the radius might be: r = 1 minute – adjacent room, r = 2 minutes – location in adjacent floor level, r = 5 minutes – location anywhere in the building, r = 10 minutes – location that is 2 buildings away, and r = 15 minutes – location that is 3 buildings away. Note that r is measured as walking time in minutes, so r = 5 implies that the real location of the user trajectory is substituted by any random location within the same building. In this example, assume, the user's visit to classroom X of building A is highly probable (i.e., he attends a class regularly on the same day of the week throughout the semester), then with the proposed technique with r = 15 minutes, classroom X will be replaced by classroom P of building D and the masked trajectory $\acute{T}_i$ will be <*library, **classroom P of build. D**, food court, classroom Y of build. B*>. The perturbing mechanism is illustrated in

Fig. 1. To further enhance privacy, certain proportion of the original trajectory (i.e., 25%, 50% etc.) can be masked using the same technique described above.

Privacy Loss Quantification: First, we need a way to quantify the information loss for a user whose real trajectory is exposed. For this purpose, we utilise entropy, a well-known approach to measure unpredictability. Entropy measures the information in a given probability distribution. To compute the entropy for trajectory T_i (in 3-hour time window), we first assign for each location $l \in T_i$ a score $score(l)$:

$$score(l) = \frac{\text{Stay time at location } l}{\text{total stay time across all locations}} \tag{1}$$

where stay time represents continuous episode of staying (i.e., not moving) at a particular location in a 3-hour time window. This score is observed across multiple time windows in a day and also across multiple days to derive probability as an average of the scores observed.

Example 2. Consider user U_1 has following trajectory:

TABLE I
USER U_1'S TRAJECTORY FOR EACH TIME WINDOW FOR 2 DAYS.

Days	Trajectories	Score
D1, TW1	$< l_1, l_2, l_3 >$	(1/3, 1/3, 1/3)
D1, TW2	$< l_1, l_3, l_4 >$	(1/3, 1/3, 1/3)
D1, TW3	$< l_1, l_2, l_5 >$	(1/3, 1/3, 1/3)
D2, TW1	$< l_2, l_4, l_6 >$	(1/3, 1/3, 1/3)
D2, TW2	$< l_3, l_5, l_6 >$	(1/3, 1/3, 1/3)
D2, TW3	$< l_4, l_5, l_6 >$	(1/3, 1/3, 1/3)

So his predicted trajectory will be $T_1 = < l_1, l_2, l_3, l_4, l_5, l_6 >$ with associated probabilities of 1/6 for each location.

We denote entropy H for the trajectory T_i as: $H(T_i) = -\sum_{l \in T_i} Pr(l) log(Pr(l))$.

Now we use the perturbation technique mentioned above (associated with the relevant obfuscation parameter), to mask the most probable locations of a user's trajectory. Once the perturbed trajectory is generated, probability for each masked location $\acute{l}$ is estimated as the global popularity of it, and estimated by:

$$Pr(\acute{l}) = \frac{\text{number of users at location } l}{\text{total number of users}}. \tag{2}$$

Entropy for perturbed trajectory $\acute{T_i}$ is then calculated as $H(\acute{T_i}) = -\sum_{\acute{l} \in \acute{T_i}} Pr(\acute{l}) log(Pr(\acute{l}))$.

We use the difference between two entropy measures $\Delta_H = H(\acute{T_i}) - H(T_i)$, to determine the entropy gain – improvement in unpredictability of user trajectories. Assume that the sensitive locations are the mostly stayed/frequent locations. In our experiments we consider obfuscating most frequent locations ($x\%$ of the locations in a trajectory).

C. Effect of Obfuscation on Performance

The twin goals of the crowd-sourcing platform are (1) to provide sufficiently accurate recommendations to the workers (i.e., minimizing the detour the worker might incur), and (2) achieve higher task completion rate, despite the inherent deliberately-introduced uncertainty in the obfuscated user trajectories.

Average Detour: Travel detour is critical in mobile crowdsourcing, as workers' physical presence at the task locations is necessary. However, when the server tries to match workers to tasks using sanitized trajectories, workers may need to travel farther to perform tasks.

For a given set of tasks, we first generate the list of recommendations (matching workers and tasks) by applying a state of the art algorithm described in [3]. To evaluate the effect of obfuscation on performance of the system, we first calculate average walking distance per user basis as detour since the user has to deviate from his original travel pattern. We denote the task location as L, location the user stayed before and after visiting the task location as X, and Y. The detour time is then $(t_{X,L} + t_{L,Y}) - t_{X,Y}$, where $t_{X,L}$ denotes the travel time to reach location L from location X.

This metric is compared against the average detour incurred when the sanitized trajectories are used to generate recommendations. Note that in the second scenario, the original (true) traces of users' trajectories are used to calculate the actual detour from recommended tasks. The difference between these two average measures helps us to understand the quality of the recommendations generated when the locations are perturbed. This can be one of the direct measures of performance loss.

Worker Productivity: To measure a worker's productivity, we calculate the average number of tasks he can complete in a time window under certain detour budget for both (a) non-private (no obfuscation is done to the user trajectories), and (b) private (where certain proportion of the trajectory is obfuscated) scenarios. The intuition is as we obfuscate the trajectories, the users may need to deviate significantly from their routine path to perform the specified tasks. This may result in a user completing a smaller number of tasks, thereby reducing the rewards earned. This earning loss may be viewed as the "privacy penalty".

V. EXPERIMENTAL EVALUATION

In this section we present experimental evaluation of our method. We utilised the Wi-Fi indoor localisation data to derive user trajectories and the state of the art algorithm described in [3] to generate task recommendations. Tasks are evenly distributed among 5 academic buildings plus the concourse of SMU campus (600 tasks per building per day) while allowing each task to be completed by 3 users. We consider 750 students who registered with the *TA\$Ker*App during our trial period September 19th – October 28th.

Privacy Evaluation Fig. 2 shows how we address the identification/profiling threats by perturbing the longer-stay locations (sensitive locations) – we present the entropy gain (in Y-axis) by perturbing (with various obfuscation parameters , depicted in X-axis) certain proportion of the trajectories. We see that, as expected, the gain in entropy/uncertainty increases as we obfuscate the sensitive location with a location that is farther away from it (i.e., higher *r*). However we also notice that the gain is not very significant when we obfuscate further – the difference in gain when *r* = 5 minutes vs. 15 minutes is nearly negligible. This suggests that even a modest degree of obfuscation *r* = 5 (building-level randomization) can deliver a decent entropy gain (around 53%), with only 25% of the trajectories are masked. Further, a similar trend is observed while varying the proportion of the individual trajectories masked (in our experiment we masked 25%, 50% and 75% of individual trajectories)– the gain in entropy increases as we obfuscate more locations of the trajectory. This suggests that even a simple, but targeted, perturbation of a user's trajectory can effectively increase the uncertainty.

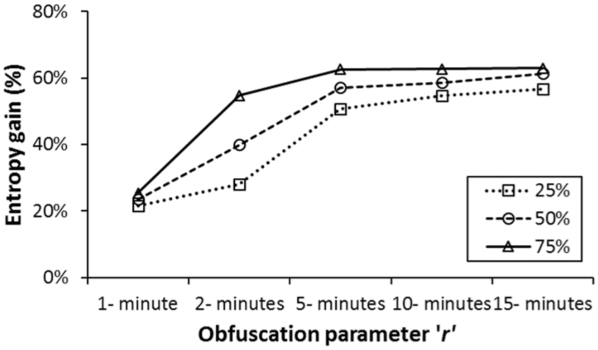

Fig. 2. Entropy Gain

Performance of the System In Fig. 3(a) we depict the performance loss (difference in worker travel distance when private and non-private mechanisms are used) vs. various *r* values. We note that, for smaller *r* values (≤ 5 minutes), the privacy-imposed obfuscation does not significantly affect the worker travel distance. This increase is, however, significant at higher values of *r*. We see that with *r*= 5, a user will incur 48% more detour (1.8 minutes more, compared with his true average detour of 3.8 minutes) while 75% of his trajectories are masked. At this operating point, the system effectively provides 62.5% more uncertainty (higher entropy) in user trajectories, while increasing the average detour experienced by ≈50%. (The 50% increase translates to a relatively modest 1.8 minute increase in absolute terms, which is not very significant in the context of SMU's urban campus).

In Fig. 3(b) we illustrate the average number of tasks a user can complete with a detour budget of 30 minutes. On Y-axis, we plot the percentage of recommended tasks completed when trajectories are perturbed as compared to the non-private system. In the hypothetical private scenario, we assumed users complete the tasks in the order of recommendation. We see that with the increasing *r*, the users complete lesser number of tasks compared to the system where such perturbations are absent. Moreover, as more locations in the trajectories are obfuscated, the completion rate drops even further, specially for larger values of *r*= 10 and 15 minutes. Specifically, with 75% of locations obfuscated, a user will gain 62.5% more uncertainty (privacy), while sacrificing detour (50% more – equivalent to 1.8 minutes) and rewards (35% lesser rewards – equivalent to $1.10). This gives us hope that by carefully choosing the obfuscation parameter, the user can achieve greater level of privacy without incurring prohibitive increase in detour or loss in rewards.

VI. Discussion

Impact on Pricing: As we see, imposing privacy techniques result in longer detours than expected. One way to compensate the worker for this additional travel overhead will be to assign higher rewards to the individual tasks. However, it will have a negative impact on the task requester as he has to invest more to ensure that the task remains attractive to the users.

Composable Privacy: In this proposed mechanism, we only considered how to protect location privacy for each individual round of recommendation (each 3-hour window) in isolation. However when the system has to generate recommendations across multiple days/weeks, users may generate traces and reveal their *new* sensitive locations where they are forced to perform tasks. Therefore, it is important for us to study the possible gradual release of private information over time.

Generalizing our Study: We note that this experiment was carried on an urban campus (4 academic buildings and 1 library – indoor campus), and additional measures should be taken to study the feasibility of applying the same technique on large scale city-wide deployments. In our prior studies [3], [4], we have shown that other crowd-sourcing related characteristics (such as analysis of worker supply, worker behaviours etc.) are consistent with observations made for city-wide crowd-sourcing. Whether this consistency applies to privacy-related observations is an open question.

Location Privacy: In this work, we aimed to protect the privacy of sensitive locations in the trajectory, and we do not aim at providing privacy for the entire trajectory (i.e., in this paper we don't consider the temporal aspect of locations). More specifically we reserve for future work questions such as: (1) how unique is a user's trajectory before and after obfuscation? and (2) how distinguishable is a temporally cloaked trajectory as opposed to the original one?

VII. Conclusions

In this paper, we showed that a simple trajectory obfuscation technique to provide location privacy, can be effective enough in providing decently accurate recommendations while enabling the participation of workers without compromising their location privacy. This trajectory obfuscation is essential to ensure that privacy is protected even after releasing the routine trajectories. As an initial exploration, we proposed a

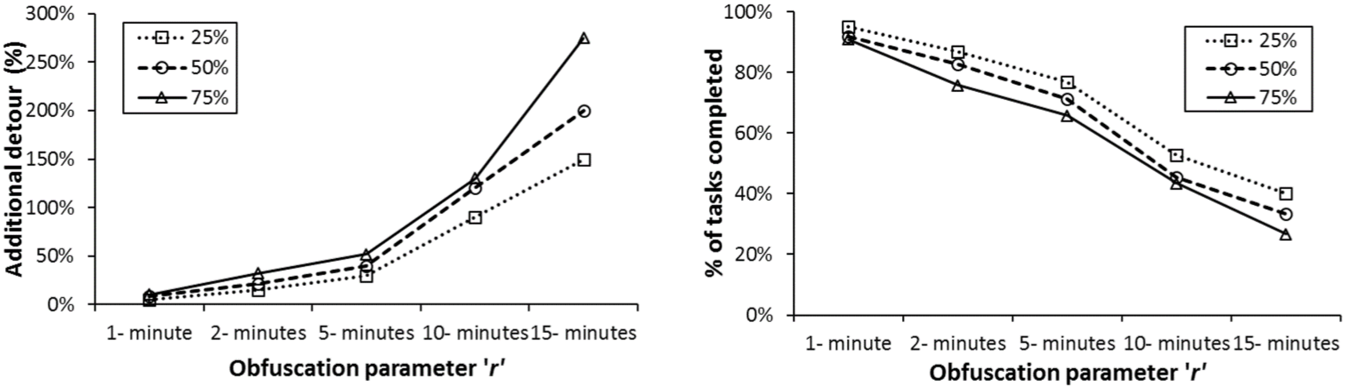

(a) Detour

(b) Productivity

Fig. 3. Efficiency of the System: (a) Additional Detour Incurred and (b) % of Tasks Completed

simple algorithm that effectively obfuscates the more dominant (higher stay time) locations of individual user trajectories, and empirically investigated obfuscation parameter settings that achieve high location privacy with low detour. Using experimental results on mobility traces data we demonstrated that the proposed technique is effective: even if 75% of the locations in an individual trajectories are perturbed, a user will end up with relatively modest increases in detour (about 1.8 minutes) and loss of rewards (approx. $1.10), while gaining 62.5% of uncertainty in his trajectory. This gives us hope that by carefully choosing the obfuscation parameter, the user can achieve high levels of privacy with relatively little loss in productivity.

ACKNOWLEDGEMENT

This research is supported by the National Research Foundation, Prime Ministers Office, Singapore under its International Research Centres in Singapore Funding Initiative.

REFERENCES

[1] Y. Cheng, X. Li, Z. Li, S. Jiang, Y. Li, J. Jia, and X. Jiang, "Aircloud: A Cloud-based Air-quality Monitoring System for Everyone." in *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*, 2014.

[2] Y. Chon, Y. Kim, and H. Cha, "Autonomous Place Naming System Using Opportunistic Crowdsensing and Knowledge From Crowdsourcing." in *12th ACM/IEEE International Conference on Information Processing in Sensor Networks*, 2013.

[3] T. Kandappu, A. Misra, S.-F. Cheng, N. Jaiman, R. Tandriyansiyah, C. Chen, H. C. Lau, D. Chander, and K. Dasgupta, "Campus-scale mobile crowd-tasking:deployment & behavioral insights," in *The 19th ACM Conference on Computer-Supported Cooperative Work and Social Computing*, 2016.

[4] T. Kandappu, N. Jaiman, R. Tandriyansiyah, A. Misra, S.-F. Cheng, C. Chen, H. C. Lau, D. Chander, and K. Dasgupta, "Tasker: Behavioral insights via campus-based experimental mobile crowd-sourcing," in *2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2016.

[5] S. Hu, L. Su, S. Li, S. Wang, C. Pan, S. Gu, T. Amin, H. Liu, S. Nath, R. R. Choudhury, and T. F. Abdelzaher, "Experiences with eNav: A Low-power Vehicular Navigation System." in *2015 ACM Conference on Ubiquitous Computing*, 2015.

[6] S. Hu, L. Su, H. Liu, H. Wang, and T. F. Abdelzaher, "SmartRoad: Smartphone-Based Crowd Sensing for Traffic Regulator Detection and Identification." *ACM Transactions on Sensor Networks*, 2015.

[7] G. Cormode, C. Procopiuc, D. Srivastava, E. Shen, and T. Yu, "Differentially Private Spatial Decompositions." in *ICDE*, 2012.

[8] M. Gruteser and D. Grunwald, "Anonymous Usage of Location Based Services Through Spatial and Temporal Cloaking." in *The First International Conference on Mobile Systems, Applications, and Services*, 2003.

[9] M. F. Mokbel, C. Y. Chow, and W. G. Aref, "The New Casper: Query Processing for Location Services without Compromising Privacy." in *The First International Conference on Mobile Systems, Applications, and Services*, 2006.

[10] B. Niu, Q. Li, X. Zhu, G. Cao, and H. Li, "Achieving k-anonymity in Privacy-Aware Location Based Services," in *IEEE INFOCOM*, 2014.

[11] S. Mascetti, L. Bertolaja, and C. Bettini, "SafeBox: Adaptable Spatio-temporal Generalization for Location Privacy Protection," *Transactions on Data Privacy*, vol. 7, no. 2, pp. 131–163, 2014.

[12] T. Higuchi, P. Martin, S. Chakraborty, and M. Srivastava, "AnonyCast: Privacy Preserving Location Distribution for Anonymous Crowd Tracking Systems," in *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2015.

[13] H. To, G. Ghinita, L. Fan, and C. Shahabi, "Differentially Private Location Protection for Worker Datasets in Spatial Crowdsourcing," *IEEE Transactions on Mobile Computing*, pp. 1–14, 2016.

[14] H. To, G. Ghinita, and C. Shahabi, "A framework for protecting worker location privacy in spatial crowdsourcing," *Very Large Data Base Endowment*, vol. 7, no. 10, pp. 919–930, 2014.

[15] L. Wang, D. Zhang, D. Yang, B. Lim, and X. Ma, "Differential Location Privacy for Sparse Mobile Crowdsensing," in *IEEE International Conference on Data Mining*, 2016.

[16] C. Chen, S.-F. Cheng, A. Gunawan, A. Misra, K. Dasgupta, and D. Chander, "TRACCS: A framework for trajectory-aware coordinated urban crowd-sourcing," in *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014.

[17] C. Chen, S.-F. Cheng, H. C. Lau, and A. Misra, "Towards city-scale mobile crowdsourcing: Task recommendations under trajectory uncertainties," in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

[18] G. Ghinita, P. Kalnis, A. Khoshgozaran, C. Shahabi, and K. L. Tan, "Private Queries in Location Based Services: Anonymizers Are Not Necessary." in *SIGMOD*, 2008.

[19] L. Kazemi and C. Shahabi, "A privacy-aware framework for participatory sensing," *ACM SIGKDD Explorations Newsletter*, vol. 13, no. 1, pp. 43–51, 2011.

[20] K. L. Huang, S. S. Kanhere, and W. Hu, "Towards Privacy-sensitive Particpatory Sensing." in *Pervasive Computing and Communications*, 2009.

[21] H. To, C. Shahabi, and L. kazemi, "PrivGeoCrowd: A Toolbox for Studying Private Spatial Crowdsourcing," in *IEEE International Conference on Data Engineering*, 2015.

[22] A. J. Khan, V. Ranjan, T.-T. Luong, R. K. Balan, and A. Misra, "Experiences with performance tradeoffs in practical, continuous indoor localization." in *IEEE WOWMOM*, 2013.