



Trading Optimality for Performance in Location Privacy

Konstantinos Chatzikokolakis, Serge Haddad, Ali Kassem, Catuscia
Palamidessi

► **To cite this version:**

Konstantinos Chatzikokolakis, Serge Haddad, Ali Kassem, Catuscia Palamidessi. Trading Optimality for Performance in Location Privacy. 11th EAI International Conference on Performance Evaluation Methodologies and Tools, Dec 2017, Venice, Italy. hal-01678256

HAL Id: hal-01678256

<https://hal.inria.fr/hal-01678256>

Submitted on 10 Jan 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Trading Optimality for Performance in Location Privacy

Extended Abstract

Konstantinos Chatzikokolakis
CNRS and Ecole Polytechnique

Serge Haddad
ENS Cachan

Ali Kassem
INRIA and Ecole Polytechnique

Catuscia Palamidessi
INRIA and Ecole Polytechnique

1 INTRODUCTION

Location-Based Services (LBSs) provide invaluable assistance in our everyday activities, however they also pose serious threats to our privacy. Location data can, in fact, expose sensitive aspects of the user's private life, see for instance [4]. There is, therefore, a growing interest in the development of mechanisms to protect location privacy during the use of LBSs. Most of the approaches in the literature are based on perturbing the user's location, see, for instance, [1, 2, 5, 6]. Obviously, the perturbation must be done with care, in order to preserve the utility of the service.

Nowadays, the most popular methods (including all those mentioned above) are probabilistic, in the sense that the perturbation is done by adding noise according to some probability distribution. Indeed, it is generally recognized that probabilistic mechanisms offer a better trade-off between privacy and utility. In this abstract we focus on the approach proposed in [6], which achieves an optimal trade-off by using linear optimization techniques. The idea is to express the desired level of privacy in the form of linear constraints, and the utility as the objective (linear) function to optimize.¹ The variables of the linear program are the conditional probabilities of reporting a location y when the real one is x , and their values, once computed, completely define the mechanism.

We consider the notion of privacy proposed in [1], called *geo-indistinguishability*. A mechanism provides geo-indistinguishability if the probability of reporting a location y when the real location is x is "almost the same" as that of every other location x' at a distance $d(x, x')$ from x , where "almost the same" means that the ratio of the probabilities is bound by $\exp(\varepsilon \cdot d(x, x'))$, with ε being the level of privacy we want to obtain per unit of distance. Formally:

$$P(y | x) \leq \exp(\varepsilon \cdot d(x, x')) \cdot P(y | x') \quad (1)$$

Intuitively, this means that x is " $\varepsilon \cdot \ell$ -indistinguishable" from the other locations x' which are at distance at most ℓ from x , where ε represents the level of indistinguishability that we want to achieve per unit distance. As explained in [1], geo-indistinguishability is based on (an extended form of) *differential privacy* [3], and it inherits its appealing properties. Notably, the robustness with respect

¹In [6] the authors fix the utility and optimize privacy. We do the reverse as our notion of privacy can only be expressed as a set of constraints, not as an objective function.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

VALUETOOLS'17, December 2017, Venezia, Italy

© 2017 Copyright held by the owner/author(s).

to composition, the independence from the prior, and a natural interpretation in terms of Bayes adversaries.

For the utility loss \mathcal{U} we use a rather general notion, namely the expected distance between the real location and the reported location. This is a function of the prior distribution on the locations π , and of the conditional probabilities that determine the mechanism:

$$\mathcal{U}(P, \pi) = \sum_{x, y} \pi(x) \cdot P(y | x) \cdot d(x, y) \quad (2)$$

As explained above, the optimal values $P(y | x)$ can be determined by solving a linear program with constraints (1) and objective function (2). Unfortunately, the number of the constraints (1) is $O(n^3)$, where n is the number of locations. Hence, due to the complexity of linear programming, the method is unfeasible even when n is relatively small. To get an idea of the dimensions, consider the Quartier Latin in Paris, which has an area of about $1.5 \times 1.5 \text{ km}^2$. If we set the size of the locations to be $100 \times 100 \text{ m}^2$, we need a grid of $15 \times 15 = 225$ cells to cover the area, which means 225^3 constraints! Reducing the granularity of the grid (i.e., considering larger cells) is not a solution, because it degrades the meaning of the utility in (2), as discussed in [2].

2 REDUCING THE SET OF CONSTRAINTS

We now propose a method to reduce the number of constraints of the linear program to $O(n^2)$, thus making the application of the method feasible for typical cases like the above one. This will be at the price of some utility loss, i.e., our method will only approximate the optimal solution. We will see, however, that the loss is quite acceptable, while the gain in performance is significant.

Let \mathcal{X} be the set of locations. Let $x_0, x_k \in \mathcal{X}$, and consider a path $x_1, \dots, x_{k-1} \in \mathcal{X}$ from x_0 and x_k . Let δ be the smallest number such that

$$\sum_0^{k-1} d(x_i, x_{i+1}) \leq \delta \cdot d(x_0, x_k) \quad (3)$$

Note that in general $\delta \geq 1$ because of the triangular inequality. It is easy to see that the constraint

$$P(y | x_0) \leq \exp(\varepsilon \cdot d(x_0, x_k)) \cdot P(y | x_k) \quad (4)$$

is a consequence of all constraints of the form

$$P(y | x_i) \leq \exp(\varepsilon \cdot d(x_i, x_{i+1})/\delta) \cdot P(y | x_{i+1}) \quad (5)$$

for $i = 0, \dots, k-1$. Therefore, it is sufficient to consider a set of constraints C of the form (5), containing enough elements so to deduce all original constraints of the form (1). Namely, it is sufficient ensure that, for every $x_0, x_h \in \mathcal{X}$, there are $x_1, \dots, x_{h-1} \in \mathcal{X}$ such that all constraints (5) are in C , for $i = 0, \dots, h-1$. Then, to achieve

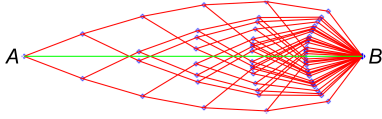


Figure 1: Possible paths between two locations A and B.

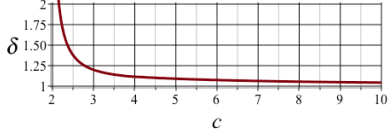


Figure 2: Worst-case value of δ as a function of c .

the original level ε of indistinguishability (i.e., of location privacy), it is sufficient to solve a new linear program with the same objective function and set of constraints C . Note that in general the solution of the new program will give an utility inferior to the original one, because the constraints in C are stricter (i.e., enforce more privacy, due to the division by δ) than the original constraints.

We construct C as follows: For every $x \in \mathcal{X}$, we consider all $x' \in \mathcal{X}$ such that $d(x, x') \leq R$, where R is some fixed distance. Then, for every $y \in \mathcal{X}$, we add a constraint of the form (5), where $x_i = x$ and $x_{i+1} = x'$. To make sure that we have enough elements in C , we assume the following *density hypothesis*: Let $ch(\mathcal{X})$ be the *convex hull* of \mathcal{X} (represented as points in the map, for instance, the centers of the cells). Then:

$$\forall y \in ch(\mathcal{X}) \exists x \in \mathcal{X} : d(y, x) \leq \rho \quad (6)$$

where ρ is some fixed distance. Note that to connect every pair of points in \mathcal{X} it is necessary and sufficient to have $R \geq 2\rho$.

The execution time depends on the cardinality of C , which, for fixed R , is $O(n^2)$ (n being the number of locations in \mathcal{X}). The cardinality of C is also proportional to R^2 , so from the point of view of efficiency it is convenient to keep R as small as possible. On the other hand, the utility is monotonic on R , hence the choice of R must take into account the trade-off between efficiency and utility.

The utility loss depends on the expansion factor δ introduced in (2), which in turn depends on R and ρ . The analysis of the worst-case for δ can be done by the geometrical construction illustrated in Figure 1. Given two locations A and B , the lines in red represent the possible paths between A and B guaranteed by the density hypothesis (6). Figure 2 shows the graph of the worst-case δ as a function of the ratio $c = R/\rho$. (Due to the condition $R \geq 2\rho$, δ is not defined for $c < 2$.) We note that δ becomes very high when c is close to 2, but it approximates rapidly the ideal value 1 as c grows.

3 EVALUATION

In this section we evaluate our method and compare with the optimal approach. We consider a set of locations disposed along the intersection points of a grid, and we set the distance u between two adjacent locations as the unit distance, i.e., all distances will be expressed in terms of u . The results illustrated in this section are valid for any value of u . For the example of the Quartier Latin, for instance, we could consider $u = 100$ m.

We note that in such grid of locations, $\rho = 1/\sqrt{2}$. Indeed, the points at maximum distance from any location are the centers of

the cells, which are at distance $1/\sqrt{2}$ from the corners of their cell. Concerning the prior, we consider a uniform distribution. We also fix $\varepsilon = \ln 2/2$, which means a level of indistinguishability 2 in a radius of $2u$. For instance, for $u = 100$ m, a user would have protection 2 in a radius of 200 m, i.e., from the point of view of an adversary, the user's real location could be no more than twice more likely than any location within 200 m from it.

We experimented with grids from $8 \times 8 = 64$ up to $15 \times 15 = 225$ locations, and values of c from 2.8 ($R = 1.98$) to 4.2 ($R = 2.97$), using an Intel machine (no TSX) 2VCPUs 2.3 GHz, 4GB RAM. The resulting computation times and utilities are shown in Figure 3. We did not evaluate the performance of the optimal method for more than 169 locations because it was taking too much time: with 169 locations it took 3,275 minutes (more than 2 days), and with 196 locations it was still running after several days.

We can see that with 225 locations like in the example of the Quartier Latin, the optimal method would be completely unfeasible. With our method and $c = 2.8$ it takes 226 minutes (note that this computation is just to *build* the mechanism, so it is done only once; afterwards, the use of the mechanism is immediate), and the utility loss is not much higher than that of the optimal method. For instance, on 169 locations, they are 3.77 and 3.49, respectively.

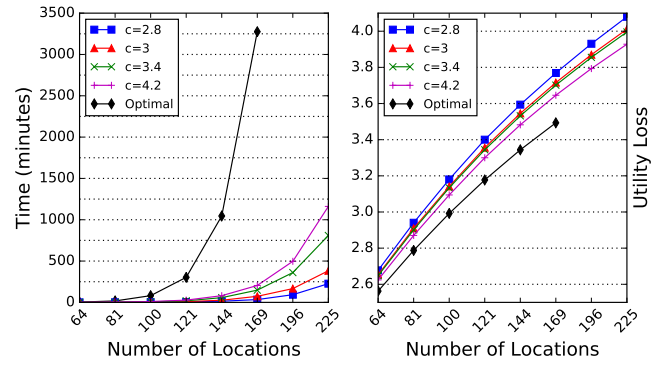


Figure 3: Execution time and utility loss.

4 FUTURE WORK

As future work, we plan to improve the bound on δ , use a prior based on real location data, and compare our method also with the one based on Laplacian noise and remapping proposed in [2].

REFERENCES

- [1] Miguel E. Andrés, Nicolás E. Bordenabe, Konstantinos Chatzicokolakis, and Catuscia Palamidessi. 2013. Geo-indistinguishability: differential privacy for location-based systems. In *Proc. of CCS*. ACM, 901–914.
- [2] Konstantinos Chatzicokolakis, Ehab ElSalamouny, and Catuscia Palamidessi. 2017. Efficient Utility Improvement for Location Privacy. *PoPETs*, 210–231.
- [3] Cynthia Dwork, Frank Mcsherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Proc. of TCC*. Springer, 265–284.
- [4] Julien Freudiger, Reza Shokri, and Jean-Pierre Hubaux. 2011. Evaluating the Privacy Risk of Location-Based Services. In *Proc. of FC'11*. Springer, 31–46.
- [5] Simon Oya, Carmela Troncoso, and Fernando Pérez-González. 2017. On the design of optimal location privacy-preserving mechanisms. *CoRR* abs/1705.08779.
- [6] Reza Shokri, George Theodorakopoulos, and Carmela Troncoso. 2017. Privacy Games Along Location Traces: A Game-Theoretic Framework for Optimizing Location Privacy. *ACM Trans. on Privacy and Security* 19, 4 (2017), 11:1–11:31.