Privacy-preserving Publication of Mobility Data with High Utility

Vincent Primault, Sonia Ben Mokhtar, Lionel Brunie
Université de Lyon, CNRS
INSA-Lyon, LIRIS, UMR5205, F-69621, France
{vincent.primault,sonia.ben-mokhtar,lionel.brunie}@liris.cnrs.fr

***Abstract*—An increasing amount of mobility data is being collected every day by different means, e.g., by mobile phone operators. This data is sometimes published after the application of simple anonymization techniques, which might lead to severe privacy threats. We propose in this paper a new solution whose novelty is two-fold. Firstly, we introduce an algorithm designed to hide places where a user stops during her journey (namely points of interest), by enforcing a constant speed along her trajectory. Secondly, we leverage places where users meet to take a chance to swap their trajectories and therefore confuse an attacker.**

***Keywords*-location privacy; data publication; time distortion; trajectories swapping**

I. Introduction

The widespread adoption of location-aware devices such as smartphones has dramatically increased the quantity of mobility data that is being continuously collected. However, collecting and sharing mobility data raises serious privacy concerns. Among the known threats is the extraction of users'*points of interest* (POIs) [1], which can be defined as places where individuals regularly spend some time, e.g., home, work, a cinema or a mall. By studying the semantics of these places, it is possible to infer sensitive knowledge like religious or political preferences. Learning users' POIs can ultimately lead to learn about the real identity of individuals with a good accuracy. Nevertheless, mobility data is still very valuable. Publishing such information allows researchers to perform real-time traffic predictions, find out interesting patterns or discover social tendencies. It is still an open and challenging issue to publish mobility traces of a set of users in a privacy-preserving manner. To reach this objective, a classical solution that is applied in the literature is to alter user's geographical locations (e.g., [2], [3]). However, this also alters the utility of published data, as trajectories are heavily distorted. This is why we propose a new solution for privacy-preserving mobility data publishing that hides users' POIs. Our challenge is to minimize the distortion of the geographical information contained in the published mobility traces. To reach this objective, our solution distorts time and opportunistically swap trajectories of users when it is possible.

The remaining of this paper is structured as follows. We present related work in Section II, an overview of our solution in Section III, before concluding in Section IV.

II. Related work

In 2002, Sweeney introduced the concept of k-anonymity. The main idea behind it is that, for each quasi-identifier of a published dataset, there must be at least k persons with this same quasi-identifier. Abul et al. proposed *Wait For Me* [3], which enforces (k, δ)-anonymity for spatial data. Their goal is to guarantee that at every instant there is at least k users at a maximum distance δ of the others. Their solution was shown to perform well with a synthetic dataset but having more difficulties to maintain a correct utility with a real-life dataset.

Differential privacy is a more recent concept introduced by Dwork in 2006. The basic idea is that an aggregate result over a database should be almost the same whether or not a single row is present inside the database. Andres et al. extended differential privacy through the concept of *geo-indistinguishability* [2], which is designed to protect location data. Like differential privacy, its strength is to provide formal and provable privacy guarantees. However, it is not yet suitable to protect mobility datasets. We have shown in [4] that, on a real-life dataset, it does not prevent the extraction of at least 60 % of the POIs even with a high privacy level.

Hoh et al. [5] proposed a solution to defeat multi-target tracking attacks. These attacks are used against mobility data protected by removing the user identifier to link back together data that is likely to belong to a same user. The proposed solution is to force paths of users to cross in areas where users are close, thus introducing confusion for an attacker. The solution has only been tested against randomized movement models.

III. Overview of our solution

The first privacy threat we want to address is the extraction of users' POIs. These correspond to places where users stop and spend some time, before moving to another place. A trace is basically a list of POIs, that can be viewed as clusters of points (as shown on Fig. 1a), with transitions in between. To hide users' POIs, we propose to enforce a constant speed along the mobility trace of each user. More specifically, we transform mobility traces to enforce an equal duration and distance between two consecutive points. If we can guarantee that speed is constant throughout a trace, it becomes difficult for an adversary to spot where a user

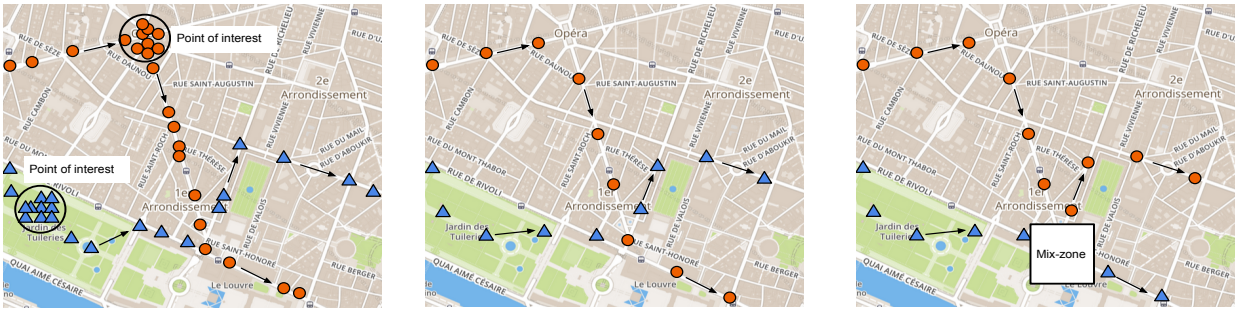

(a) Original trace (b) After enforcing a constant speed (c) After swapping trajectories

Figure 1: An example of two mobility traces anonymized with our solution.

stopped because there is no point at which she appears to be stationary. Clues can still be obtained from background knowledge (e.g. the probability is higher to stop in a park that in the middle of a motorway) but there will be no certainty for an attacker (e.g. a user can either have just crossed the park or had a picnic there). Figure 1b shows the result of this operation applied on two mobility traces. From this figure, we can see that the POIs of the users have been hidden and that points on each trace are evenly distributed.

The other privacy threat we want to address in this paper is the re-identification of users. To reach this objective, we exploit natural paths crossings to confuse an attacker. When users move during a day, they continuously meet other users in public transportations, malls, work places, etc. These meeting areas are called mix-zones, as introduced in [6]. Mix-zones are well-delimited areas where nobody is tracked; we only know where and when users enter and leave a mix-zone, without any insight about what happens inside. Before users leave a mix-zone, their identifiers are possibly shuffled. A user entering a mix-zone labelled as "user A" could either leave it labelled as "user B" or remain "user A". It therefore helps breaking the correlation between traces before and after the mix-zone. We do not want to artificially distort traces to force users to meet, but instead we take advantage of existing mix-zones. Figure 1c shows that the two traces have been swapped inside the mix-zone.

Our main utility goal was to minimally distort the location. The first step introduces only error when interpolating new points between known ones. If the sampling rate is high enough, this interpolation should be precise enough to introduce almost no spatial inaccuracy. Similarly, the second step only swap user identifiers but does not alter the location. The only utility loss comes from the fact we suppress points inside mix-zones, but this should be a reasonable degradation as long as mix-zones remain reasonably small. Tough, we acknowledge not all queries can be implemented with our solution. For example, studying transitions between locations cannot be done because users can possibly be swapped. We believe there is not one solution that fits all use cases, and that the appropriate solution must be chosen according to both users' preferences in terms of privacy and the analysts' objectives.

IV. Conclusion

In this paper we presented an overview of a new solution to anonymize mobility datasets. Its novelty resides in the fact that it obfuscates time instead of obfuscating location, which allows to have a better spatial accuracy than state of the art solutions. Furthermore, our solution swaps trajectories in order to confuse the attacker, without compromising utility. We now plan to evaluate our solution with real-life datasets and study its resiliency against common privacy attacks and the practical utility it can preserve.

Acknowledgment

This work was supported by the LABEX IMU (ANR-10-LABX-0088) of Université de Lyon, within the program "Investissements d'Avenir" (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR).

References

[1] S. Gambs, M.-O. Killijian, and M. N. del Prado Cortez, "Show Me How You Move and I Will Tell You Who You Are," *Transactions on Data Privacy*, vol. 4, no. 2, pp. 103–126, Aug. 2011.

[2] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-indistinguishability: Differential Privacy for Location-based Systems," in *Proceedings of CCS'13*. ACM, 2013, pp. 901–914.

[3] O. Abul, F. Bonchi, and M. Nanni, "Anonymization of moving objects databases by clustering and perturbation," *Information Systems*, vol. 35, no. 8, pp. 884–910, 2010.

[4] V. Primault, S. Ben Mokhtar, C. Lauradoux, and L. Brunie, "Differentially Private Location Privacy in Practice," in *Proceedings of MOST'14*, May 2014.

[5] B. Hoh and M. Gruteser, "Protecting location privacy through path confusion," in *Proceedings of SECURECOMM'05*. IEEE Computer Society, 2005, pp. 194–205.

[6] A. Beresford and F. Stajano, "Location privacy in pervasive computing," *Pervasive Computing, IEEE*, vol. 2, no. 1, pp. 46–55, Jan. 2003.