Semantic and Sensitivity Aware Location Privacy Protection for the Internet of Things

Berker Agir, Jean-Paul Calbimonte and Karl Aberer

Faculty of Computer Science and Communication Systems, EPFL, Switzerland.
firstname.lastname@epfl.ch

Abstract. Everyday applications and ubiquitous devices contribute data to the Internet of Things, oftentimes including sensitive information of people. This opens new challenges for protecting users' data from adversaries, who can perform different types of attacks using combinations of private and publicly available information. In this paper, we discuss some of the main challenges, especially regarding location-privacy, and a general approach for adaptively protecting this type of data. This approach considers the semantics of the user location, as well as the user's sensitivity preferences, and also builds an adversary model for estimating privacy levels.

1 Introduction

The Web is continuously evolving and integrating into its core thousands of smart devices, mobile phones and interconnected *things* that are capable of sensing and capturing information dynamically every second and report it to the cloud. Many challenges arise from this emerging Internet of Things, including the protection of the sensitive information of people, which can be mined and exploited by an adversary. More concretely, we focus on protection of location data, which is commonly tracked by mobile sensing applications and also applications for smartphones and similar devices that feed the Internet of Things (IoT). Location is a context-rich piece of information that is both sensitive and necessary for many IoT scenarios such as participatory sensing or location-based online services. For instance, several initiatives exist that try to crowd-sense data such as air quality in urban spaces, from both stationary and mobile devices managed by common people (e.g., AirQualityEgg[1], OpenIoT[2] or OpenSense[3]). In these cases, simple anonymization or aggregation techniques are not enough to prevent inferring private information about people contributing to the system [4], regarding their daily habits, leisure patterns or even political or religious orientations.

Most of the existing work on location-privacy approaches the problem only from a geographical perspective and assumes that if the actual location of a user is obfuscated by a region or hidden completely, it is protected. However, these techniques generally do not take into account the semantics of a user's location or his different sensitivity levels w.r.t. a certain place or type. For example, for

[1] http://airqualityegg.com/
[2] http://openiot.eu
[3] http://www.nano-tera.ch/projects/423.php

the majority of the population in a city, hospitals might be sensitive locations, because disclosing the fact that a person is visiting a hospital may reveal that he has a disease. Although in general this could mean that this location information should be hidden, for other users (e.g., a physician who works at the hospital) this might not be sensitive at all. In this particular example, we see that the sensitivity depends on the *semantics* of the location and the user's preferences. A potential adversary can take advantage of this type of information, even more, considering that this type of semantic location information can be easily and publicly accessed through different geo-tagging, and geo-Linked Data interfaces [2]. Moreover, the protection strategies should adapt to the user context, e.g. trajectories, time of day, density of people in a certain place, etc.

In this paper we propose a general approach for location-privacy protection that takes into account not only the geographical but also the semantic information of urban locations, as well as user's sensitivities to obfuscate the location information that is transmitted to a service provider of Internet of Things. This approach extends and leverages on our previous set of protection techniques [1], fundamentally adding the combination of semantics and attack models. In our approach each individual is able to build a model of an adversary, based on possible strategies and background information, so that he can adaptively change the obfuscation level according to a required level of protection.

2 Adaptive and Sensitivity-aware Privacy Protection

In this work, we focus on scenarios where users with mobile devices move in an urban area and they continuously provide data to the IoT for a specific purpose. We assume that users provide their geographical location information along other domain-specific data. Users occasionally visit certain semantically-annotated places (such as "Restaurant" or "Hospital"). In this case, the corresponding server may know what type of location is visited thanks to availability of rich online resources (geo-tagged maps, geo-Linked Data, etc.).

In this setting, location privacy of users is threatened by the server they communicate with (or any other entity who has access to the data they send). In particular, users may not want to disclose their actual locations in the fear that an adversary can exploit this information to gain more private data about them. Therefore, they would like to obfuscate (*i.e.,* reduce the granularity of) their locations or hide them completely. There have been numerous proposals for protecting location-privacy in different ways (see Section 3), but none can continuously adapt the protection strategy considering both location semantics and varying user sensitivities. In summary, we focus on two main points:

1. We model the adversary that takes into account location semantics and user sensitivities in addition to users' geographical mobility behaviors. The adversary can develop sophisticated algorithms (e.g., using machine learning techniques) for inferring users' actual locations.
2. We develop an *adaptive* user-side protection mechanism that considers location semantics and users' varying privacy sensitivities. Previous work [1]

proves that adaptive approaches retain location privacy better than static obfuscation approaches. Distinctively, our scheme automatically determines the required privacy levels based on location semantics and user sensitivities and meets these levels by anticipating the adversary's strategy.

A crucial advantage of our approach is the ability to model the adversary, and thus the possibility to continuously simulate his attacks and evaluate his probability of success. The adversary model considers the parameters that users already know: user mobility, location semantics and privacy sensitivities. A user sensitivity profile can be crowdsourced or inferred [8] and then personalized as needed. Figure 1 depicts the standing points of the adversary and users. These two problems can be seen as two sides of the same coin, in the sense that both users and the adversary reason about their counterparts' knowledge and capabilities while taking action. We explain the details of these points in sections 2.2 and 2.3, respectively, after introducing the framework formalization in Section 2.1.

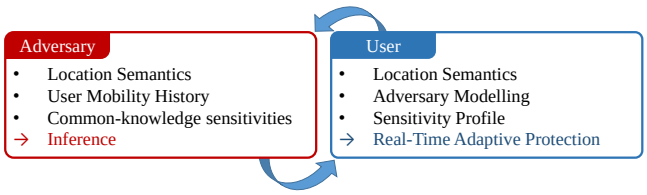

Fig. 1: Adversary and users move in opposite directions: the adversary obtains both public and side information about users in order to breach privacy; conversely, users anticipate the adversary's capabilities in order to protect their privacy adaptively and in real-time.

2.1 Formalization

As previously stated, we consider mobile users $\mathcal{U} = \{u_1, u_2, ..., u_N\}$, who move in a discretized urban area consisting of M regions in set $\mathcal{R}$. A user u generates an actual trace a_u as he moves through time period $\mathcal{T} = \{t_1, t_2, ..., t_T\}$. Each element, *i.e.,* event, in a_u is of the form $a_u(t) = \langle u, t, r, c \rangle$, where $r \in \mathcal{R}$ is user u's actual location at time $t \in \mathcal{T}$ and $c \in \mathcal{C}$ is a semantic tag for region r.

User u would like to connect to a server of a certain online application, but is afraid that his privacy will be violated. Hence, he employs a location-privacy protection mechanism f, which generates an obfuscated trace o_u from a_u. Note that o_u is the disclosed location trace of user u and it is observed by the adversary. Users may have different privacy-sensitivity levels for different semantic tags and even for some particular regions. As a result, a protection mechanism f takes into account user u's sensitivities denoted by $\mathcal{S}_u \in \mathcal{S}$.

2.2 Adversary Model

A server that aggregates data from users U in an IoT environment, is interested in inferring their private data through location context; hence he is considered as the adversary. We assume that he can obtain some background information about a user u which he exploits in order to enhance his attack. This background information is in the form of transition counts between regions, which constitutes

user u's behavior history, denoted by H_u. We argue that the adversary may fail to obtain a complete history profile H_u and try actually to approximate it. We denote by $\widehat{H_u}$ the history profile as observed and obtained by the adversary.

The adversary knows which protection mechanism f users employ and also knows that the users may provide f with their sensitivity profiles $\mathcal{S}$ to meet their privacy requirements. Most users may share similar sensitivity values for many semantic tags or at least it is possible to predict such sensitivity values, with some challenge for specific locations [8]. Therefore, the adversary may not have the actual sensitivity profiles $\mathcal{S}$ of users, but a close approximation for most of the population. In this regard, he can build a common-knowledge sensitivity profile set $\widehat{\mathcal{S}}$. Formally, this translates into the fact that, a mechanism f generates an obfuscated trace o_u from a_u according to a probability distribution $\mathbf{Pr}\{o_u|a_u, \mathcal{S}\}$. As a result, the adversary builds his attack in terms of the probability that a trace a'_u is the actual trace of user u given the adversary observation o_u, his background knowledge $\widehat{H_u}$ and sensitivity profile $\widehat{\mathcal{S}_u}$: $h(a'_u) = \mathbf{Pr}\{a'_u|o_u, \widehat{H_u}, \widehat{\mathcal{S}_u}\}$. $h(a'_u)$ represents the posterior probability distribution on user traces based on the adversary's attack. Considering that the adversary has considerably more computational power than users, but still limited, we model a sophisticated attack with machine learning techniques on multidimensional data, *i.e.,* geographical location, semantic dimension, time and user privacy sensitivities. The ultimate purpose of the attack is to approximate h.

2.3 Adaptive Privacy-Protection

We know that the adversary has a certain knowledge about users (*i.e.,* $\widehat{H}$, $\widehat{S}$, f), and therefore users (and the protection mechanisms they employ) must be aware of this fact. Hence, we build our location-privacy protection scheme in an adaptive manner; it reasons about the information the adversary has, anticipates on what he can infer from disclosed data and decide on the protection details accordingly. It also regards the semantics of the user's location and his privacy sensitivities, which it integrates in its decision process. Formally, the protection mechanism f employed by a user u generates an obfuscated event $o_u(t) \in o_u$ given the actual event $a_u(t)$, all the past disclosed events $\{o_u(t')\}$ for $t_1 \leq t' < t$, user u's history profile H_u, and his sensitivity profile $\mathcal{S}_u$:

$$o_u(t) = f(a_u(t), \{o_u(t')\}, H_u, \mathcal{S}_u), \qquad t_1 \leq t' < t$$

In order for f to exploit the aforementioned data in addition to the user's actual location and also to adapt to user behavior in real-time, it needs to reason about the adversary's strategy, *i.e.,* how the adversary fuses all the data he has for inference. This means that the user should attack his own obfuscated trace and make an evaluation of the expected privacy level for his protection mechanism's actions before disclosing any data. However, users do not have as much computational power as the adversary has; hence, they need to model a more efficient and thus weaker attack in order to approximate what the adversary can achieve. For this, we introduce two separate core modules in our adaptive privacy-protection

scheme: the module responsible for local estimation of location-privacy and also the protection module that essentially apply the protection techniques on user locations. Figure 2 shows these modules with their interaction among each other.

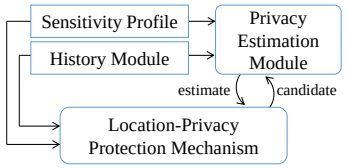

Fig. 2: We model the adaptive protection scheme based on two core modules: local privacy-estimation module and protection module. They interact with each other for adaptation in real-time by utilizing user history and sensitivity profile.

Privacy estimation module (PEM) keeps track of the user's past events and history profile. It then uses Bayesian inference to attack user's own trace as described in [1]. For efficiency, this is achieved through storing user (obfuscated) events in an inference graph and updating the graph in real-time as the user generates new events. PEM fuses information from Bayesian inference with history data and computes a privacy level using the expected distortion metric proposed by Shokri *et al.* [6]. This metric is basically an expected value computation on distances between the user's actual location and the observed locations in his obfuscated event. Note that what the user estimates here is in fact the posterior distribution h resulting from the adversary attack. However, the adversary attacks the obfuscated trace as a whole; the user only attacks the disclosed part of his obfuscated trace. As a result, the user achieves an approximation of h.

Whenever the user generates a new event, the protection module obtains the actual location and the corresponding semantic tag; then it checks the user's sensitivity profile $\mathcal{S}_u$ and drafts an obfuscated location by also considering the user's history profile H. It invokes the PEM by passing it the generated obfuscated location and the PEM evaluates the expected privacy-level as if the user would disclose the current obfuscated location. Upon receiving the estimated privacy-level, the protection module checks if it satisfies the user's sensitivities. If yes, it discloses the obfuscated location. Otherwise, it adjusts its parameters and generates a new obfuscated location, and goes through the same procedure. In summary, the protection mechanism iteratively adjusts its obfuscation parameters until the user's sensitivity preferences are satisfied.

3 Related Work

There have been numerous works in the literature on location privacy, primarily variations of protection mechanisms that make use of obfuscation, perturbation and hiding. Most of the proposed solutions are evaluated based on static parameters (e.g., obfuscation area size) which lacks a potential adversary's perspective and therefore do not yield realistic privacy evaluation. Shokri *et al.* [6, 7] addressed this problem and proposed a framework for quantifying location-privacy w.r.t. an adversary with limited capabilities, but some background knowledge on user behavior. They evaluate users' location privacy based on the result of the adversary's attack, *i.e.,* his confusion, correctness and accuracy, but they developed their framework only on geographical dimension of location.

Additionally, some prior work takes into account user sensitivities and location semantics in order to better protect location privacy of users. For instance, the PROBE Framework by Damiani *et al.* [3] propose to generate personal obfuscation areas based on users' privacy sensitivities to certain types of places and their ratio to their obfuscation areas in terms of size. Another work, by Monreale *et al.* [5], focuses on users' semantic trajectories which consist of sequences of visited places. They argue that some of the places of certain types in a trajectory might be sensitive and develop a protection approach based on semantic generalization according to a predefined semantic tag taxonomy. Both of these work analyze location-privacy by computing expected confusion levels without considering an adversary's capabilities and attack strategies.

4 Conclusions & Future Work

We presented an adaptive location-privacy protection scheme that is aware of location semantics and user sensitivities. This scheme takes into account a sophisticated adversary by emulating his attack and thus estimates users' expected privacy levels. We also model and formalize this sophisticated adversary for realizing the threat, through which we aim to evaluate location-privacy comprehensively. Even though less powerful than an actual adversary, we expect our scheme to perform better against a sophisticated adversary than previous work. Our next step is to develop both the types of attacks described and the adaptive protection scheme. Moreover, user sensitivities may depend on more features such as time of day and the user activity. We will investigate these aspects and their effect on privacy in order to enhance and complete our framework.

Acknowledgements: This work is supported by the OpenSense2 project, funded by Nano-Tera.ch and financed by the Swiss Confederation.

References

1. B. Agir, T. G. Papaioannou, R. Narendula, K. Aberer, and J.-P. Hubaux. User-side Adaptive Protection of Location Privacy in Participatory Sensing. *Geoinformatica*, 18(1):165–191, 2014.
2. S. Auer, J. Lehmann, and S. Hellmann. *Linkedgeodata: Adding a Spatial Dimension to the Web of Data*. Springer, 2009.
3. M. L. Damiani, E. Bertino, and C. Silvestri. The PROBE Framework for the Personalized Cloaking of Private Locations. *Trans. Data Privacy*, pages 123–148, 2010.
4. P. Golle and K. Partridge. On the Anonymity of Home/Work Location Pairs. In *Pervasive Computing*, volume 5538, pages 390–397. Springer, 2009.
5. A. Monreale, R. Trasarti, D. Pedreschi, C. Renso, and V. Bogorny. C-safety: A Framework for the Anonymization of Semantic Trajectories. *Trans. Data Privacy*, pages 73–101, 2011.
6. R. Shokri, J. Freudiger, M. Jadliwala, and J.-P. Hubaux. A Distortion-based Metric for Location Privacy. In *WPES*, 2009.
7. R. Shokri, G. Theodorakopoulos, J.-Y. Le Boudec, and J.-P. Hubaux. Quantifying Location Privacy. In *IEEE S&P*, pages 247–262, 2011.
8. E. Toch. Crowdsourcing Privacy Preferences in Context-aware Applications. *Personal and Ubiquitous Computing*, 18(1):129–141, 2014.