



HAL
open science

Privacy protection control for mobile apps users

Sophie Cerf, Bogdan Robu, Nicolas Marchand, Sara Bouchenak

► **To cite this version:**

Sophie Cerf, Bogdan Robu, Nicolas Marchand, Sara Bouchenak. Privacy protection control for mobile apps users. *Control Engineering Practice*, 2023, 134 (May), pp.105456. 10.1016/j.conengprac.2023.105456 . hal-03977386

HAL Id: hal-03977386

<https://hal.science/hal-03977386>

Submitted on 7 Feb 2023

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.

Privacy protection control for mobile apps users

Sophie Cerf^{a,*}, Bogdan Robu^b, Nicolas Marchand^b and Sara Bouchenak^c

^aUniv. Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRISTAL, Lille, F-59000, France

^bUniv. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, Grenoble, 38000, France

^cINSA Lyon - LIRIS - CNRS, Distributed Systems Research Group, Lyon, France

ARTICLE INFO

Keywords:

control of computing systems
location privacy
differential-privacy
modeling
Sampled-Data Control

ABSTRACT

Predominant in today society, mobile apps are rising as promising application systems for automatic control. An app can be viewed as a plant, processing input signals (queries, phone data, etc.) and generating outputs (such as a service or an answer). Guaranteeing that the app complies with a desired behavior is a major safety challenge. This work focuses on privacy issues for geolocated mobile apps. Many applications use the location data to provide a service (e.g., navigation, fitness) or to improve it (e.g., weather forecast, social media). This gain in service utility comes at the cost of personal data sharing. Such threat to user privacy can be leveraged by protection mechanisms, e.g., addition of noise to the location data. However, state-of-the-art techniques still lack means of ensuring both data utility and privacy in a dynamics utilization context. This paper presents the first non-linear analytical modeling followed by a control formulation for regulating the privacy level in a mobile app. The privacy is sensed using the well established notion of Point of Interest. Through modeling, we highlight the control challenges, namely the non-linearity and time-variance of the plant, its high sensibility to noise and the impact of the user's mobility pattern – seen a disturbance. A controller is designed, combining feedback with anticipation action. Evaluation is performed using mobility records from two real-world multi-users datasets. Our approach enables, with a unique and universal tuning, to robustly meet privacy objectives with preserved utility and negligible computational overhead. Control algorithm, experimental evaluation and analysis scripts are available online for reproducibility.

1. Introduction


Computing systems, and especially software, are novel control systems. The need for IT regulation started in the 2000s when the concept of Autonomic Computing [30] was introduced, aiming at the self-adaptation of software products without being backed by a specific theory. State-of-the-art works investigate the use of control theory for modeling and decision-making in computing systems [18], showing strong and promising opportunities. Software systems evolve with unseen dynamics, as they do not fall under Physics' laws, therefore, novel opportunities are thus open for research on all control aspects: formulation and choice of sensors and actuators, modeling through identification from data, and control [34]. Applications running on mobile devices represent a significant part of IT usage by individuals. Ensuring their behavior regulation is a major safety challenge. Moreover, these apps, dynamic by nature, offer a whole new playground for control specialists [44]. In this work, we tackle privacy challenges for users of mobile apps using geolocation.

Location privacy protection considers mobile devices users whose mobility information is shared with third parties. Applications and services tend to require location data

to personalize users' experiences. Examples of location-based services are very numerous and range from navigation applications and fitness tracking apps, to weather forecasting and recommendation systems [24]. Mobile apps provide those personalized and convenient services at the cost of personal data disclosure, as service providers or third party attackers can take advantage of these data to derive private information about users. The most common threats are (i) re-identification attacks where the identity of an anonymous user is guessed based on previously recorded data [21], (ii) mobility prediction that anticipates users' next moves based on their habits [23], (iii) extraction of user's places of interest [35] (home, workplace, place of worship, etc.), and (iv) inference of social relationships (partners, coworkers, etc.) [46]. Such severe and global threats for users motivate the need for efficient and accessible privacy protection. Additional motivation concerns political and societal aspects, as illustrated by the media release of a mobility dataset by a fitness tracking social app that revealed the location and maps of unknown US military bases [48]; or the regulations enforced by the governments, such as the European General Data Protection Regulation¹ or the US Location Privacy Protection Act².

The time-perspective plays a central role in mobile apps, and can suggest a classification [27]: some are snapshot services—they only need a single location point to provide their service, an example could be a weather app—, while the others need continuous records—for instance navigation or gaming apps. Without loss of generality, this work considers

*Corresponding author

 sophie.cerf@inria.fr (S. Cerf); bogdan.robu@gipsa-lab.fr (B.

Robu); nicolas.marchand@gipsa-lab.fr (N. Marchand);

sara.bouchenak@insa-lyon.fr (S. Bouchenak)

 <https://sites.google.com/view/sophiecerf/> (S. Cerf);

<http://www.gipsa-lab.grenoble-inp.fr/~bogdan.robu/> (B. Robu);

<http://www.gipsa-lab.grenoble-inp.fr/~nicolas.marchand/index.html> (N.

Marchand); <https://perso.liris.cnrs.fr/sara.bouchenak/> (S. Bouchenak)

ORCID(s): 0000-0003-0122-0796 (S. Cerf)

¹<https://gdpr-info.eu/>

²<https://www.congress.gov/bill/112th-congress/senate-bill/1223>

the more complex, online scenario where a stream of data is dynamically shared. As the attacker has access to more information, the continuous scenario is more challenging than the snapshot case.

Location Privacy Protection Mechanisms emerged as solutions to protect users' privacy [43]. Such algorithms modify the location data to improve privacy; e.g., by adding noise [5], reducing data precision [13], or merging close users' locations [1]. In this paper, we focus on the well-established notion of differential privacy [15] and its adaptation to mobility data, Geo-Indistinguishability [5]. This mechanism adds spatial noise to each of the location data. Most protection mechanisms are parametrized algorithms, e.g., the variance of the spatial noise added to the data can be adjusted. By tuning those parameters we change the protection action. This property is highly valuable considering that privacy often comes at the cost of a reduction of the service utility: a configurable mechanism which allows to leverage the privacy to utility trade-off [10]. In particular, optimal privacy and utility protection cannot be universal [8], motivating the need for dynamic adaptation of the protection. Deciding on the suitable parametrization to meet privacy and utility objectives at runtime is however still an open challenge. As human mobility is highly dynamic, with varying speeds and frequencies of move, the application of a protection mechanism with a fixed configuration results in volatile levels of privacy and utility.

This paper adopts a novel approach, in which the problem is shifted from choosing the configuration parameter of a protection mechanism to achieving a desired privacy level in an automated and robust fashion while improving utility. This feedback approach allows therefore a gain of utility even in non-privacy-sensitive situations. The challenge of how to automatically tune a protection mechanism to meet users objectives, is shaped as a reference tracking problem. When it comes to ensuring protection independent of the user mobility pattern, a disturbance regulation approach is used. This paper thus tackles the challenges of practicality, personalization of protection mechanisms, and temporal aspects [27] thanks to the control-based approach. This works illustrates the direct application of control theory for the benefit of users of mobile communication systems.

In this paper, we introduce a feedback control strategy for mobile privacy regulation, as illustrated in Fig. 1. In details, we present the first formulation of the location privacy challenge as a Control problem. Online privacy relies on user's likelihood to be in a point of interest; and utility is based on service quality loss. Control challenges rely on non-uniform sampling, saturation effects, and stochastic, highly dynamic and unpredictable disturbance. Modeling is performed both analytically and by data identification to derive a dynamical model with input non-linearity. To the authors' knowledge, it is the first analytical non-linear modeling of a software system. A privacy controller is designed, combining feedback and a pre-compensation. Its parameters are universal, and no user-specific tuning is needed. Evaluation on several real mobility data illustrates the relevance of the control

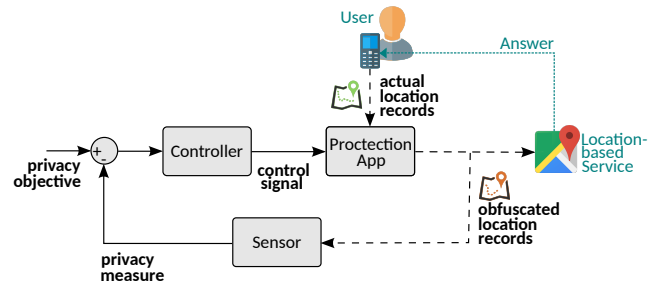


Figure 1: Privacy regulation of a mobile protection app using control: a schematic view.

approach. Control algorithm, evaluation and analysis scripts are available online for reproducibility [11].

In short, the main contributions are:

- (i) formulation of the location privacy control problem,
- (ii) non-linear analytical modeling, and
- (iii) design of a privacy controller.

The remaining of the paper is organized as follows. First background on mobility and privacy is given in Section 2, further translated into a control problem in Section 3, highlighting dynamic protection challenges. The system's model is presented in Section 4, with both analytical and identification approaches. The controller construction and design is presented in Section 5. Validation and evaluation on real data ends the paper in Section 6, followed by related work analysis (Section 7) and by the conclusion with perspectives. (Section 8).

2. Privacy protection of mobile apps users

Our control plant is a location-privacy protection application and mobility data is considered as a disturbance. This section provides background on the protection mechanisms and on mobility, while control formalization of the considered problem is given in Section 3. First, we highlight the challenges of controlling a system influenced by human mobility data: (i) wide-spectrum signal frequencies, (ii) large amplitude variations and (iii) non-uniform sampling. Then we present the real-world datasets used for identification and control evaluation. Eventually, background is given on the notions of location privacy and protection mechanisms.

2.1. Mobility data: a challenging disturbance

Mobility data are records of location coordinates over time. The location is often the one of an individual using a mobile device, but it can also be records of some transportation modes (car, bikes, etc.). In the following, the terminology *user* is used to refer to the person or object which location is recorded. Formally, mobility databases are GPS points (latitude, longitude) labeled with timestamps and users' identification. Depending on the data collection method, the sampling rate can be constant or not during the record period, and some periods can even contain no record (for instance when the device is switched off).

In this work, we apply our research on two datasets, namely: Cabspotting and Privamov. The Cabspotting dataset

[40] gathers mobility traces of 536 taxi cabs during their service in San Francisco Bay Area, USA, over one month. The sampling time is not constant, with an average of 1 min, and the dataset is very dense—few data are missing. Note that the specificity of this dataset is regarding the type of mobility it represents: only vehicles moves, mostly in a dense urban area. The Privamov dataset [7] records the mobility of students and academics around Lyon’s Campus, France. Data are recorded with a fixed sampling period of 10 s, however much data is missing which makes the dataset very sparse. This dataset represents human urban mobility with various transportation means.

A snippet of a mobility trace from a Cabspotting user can be found in Fig. 2: the location points are represented on the map of San Francisco (Fig. 2a) and the variations of the user’s speed through time are given (Fig. 2b). Some sections of this trace have been highlighted to show mobility diversity. First, the cab driver is stopped at a cab station (yellow dots, between 0 min to 40 min). Later on, the user drives with high speed on a highway from the airport to the city (blue crosses, around 110 min to 120 min). Eventually, the user drives around the city, from a hotel to the piers with relatively slow speed (green circles, around 150 min). The user mobility is extremely different with, various signal amplitude, frequency, and sampling time. One can also note that, for instance, no data is recorded from 5 min to 30 min. Robustness against such a disturbance is a real challenge.




Figure 2: Mobility data of a user from Cabspotting dataset (abooip) revealing the diversity of mobility patterns.

2.2. Privacy Protection

Before describing the protection mechanisms (with example and illustration), we focus on the thorny notion of privacy.

2.2.1. Notions of Location privacy

Defining privacy is not easy as there are flourishing attempts in the literature, mainly around the theoretical concepts of k -anonymity [47] (hiding a user among $k - 1$ other similar ones) and differential privacy [15] (bounding the query answers difference between datasets differing by only one user). Those theoretical concepts are ill-suited for practical scenarios, as the privacy level depends on contextual information. Indeed, both the semantic context (for instance religious or health related places) and user context (history

in this vicinity) impact users’ privacy threat. In practice, location quantification based on metrics are better suited than theoretic definitions. There are many practical location privacy metrics (see Primault et al. [43] for an overview), but one key concept is the notion of Points Of Interest (POI) [26, 19], which regroups the significant places where users spend time, such as home, workplace, place of worship, etc. A POI is defined as a meaningful geographical area where a user made a significant stop [26]. Formally, it is the location of the center circular zone of a given diameter where the user stayed for a minimum duration T .

The retrieval of POI is of crucial knowledge as it is often the very first step for performing attacks as re-identification [41], mobility prediction [23] or worship prediction [20]. This work does not try to defend against specific attacks, but rather focuses on ensuring user-defined levels of a privacy metric based on the well-established notion of Point of Interest [26]. Indeed, the protection challenge aiming to defeat a well-performing privacy attack has been shown unfeasible in practice, as it would require adding so much noise that the data would be unusable by the mobile app [32]. On a more theoretical level, not specific to location, the optimal state estimation of systems under Laplacian noise (i.e., differential privacy) has also been investigated [17]. Additionally, from a practical point of view, we address privacy by using a privacy sensor signal based on POIs (see Section 3.1.4). The *obfuscation* of the users’ Points Of Interest would be our notion of location privacy.

2.2.2. Protection mechanisms

Location Privacy Protection Mechanisms (LPPMs) is the literature dedicated terminology for all the processes and algorithms that, by manipulating location data, aim at improving the privacy protection of users. They take as input a mobility record and output another mobility data, hopefully more privacy preserving. The input data is called actual data, or original one; while the output is called obfuscated or sanitized data.

We want to highlight that the methodology presented in this work is general and can apply to all existing protection mechanisms [43] which are:

- (i) tunable by at least one parameter;
- (ii) online processes: every location can be individually obfuscated in real time;
- (iii) user-centric: the obfuscation does not depend on external knowledge (like: crowd density in the area, other users’ position, etc.).

In this paper, we use the Geo-Indistinguishability mechanism (Geo-I) [5] as an illustration of our control approach as it is one of the most used mechanism which realizes spatial distortion of user mobility data in an on-line fashion (i.e., suitable for applying control), in opposition with most of the state-of-the-art protection app which work in an off-line fashion. Moreover, Geo-I realizes the well known differential privacy model [15], which enable to derive mathematical privacy guarantees on the sanitized dataset. Geo-I protects user’s location data by adding spatial noise. This noise is

drawn from a Laplace distribution around the actual user location, and repeated for each record through time. The noise is stochastic to avoid easy filtering by an attacker knowing the data protection conditions. It has a configuration parameter $\epsilon \in \mathbb{R}^+$ (expressed in m^{-1}) which quantifies the amount of noise to add: the lower ϵ , the higher the noise. For an actual location point l (which is a vector of latitude and longitude coordinates), its sanitized (obfuscated) value named 1 is computed as follows:

$$1 = l - \frac{1}{\epsilon} \left[W_{-1} \left(\frac{p-1}{e} \right) + 1 \right] \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix} \quad (1)$$

where W_{-1} is the Lambert W function (the -1 branch), e is Euler's number, p is drawn uniformly in $[0, 1)$ and θ in $[0, 2\pi)$. Note that the use of the Lambert W function comes from Geo-I's formulation. Using this specific function—rather than a Gaussian distribution—ensures the well established differential-privacy property [15] on the protected data. For the sake of simplifying the notations through the rest of the paper, we introduce the following notation:

$$\mathcal{W}(p) = W_{-1} \left(\frac{p-1}{e} \right) + 1 \quad (2)$$

which rewrites Eq. (1) as follows:

$$1 = l - \frac{\mathcal{W}(p)}{\epsilon} \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix} \quad (3)$$

Fig. 3 illustrates the application of Geo-I protection app on the mobility trace of an illustrative Cabspotting user, for two values of the parameter ϵ (Figs. 3b and 3c) compared to the actual original data trace (Fig. 3a). The noisier the data are, the better the user privacy is preserved, as less information can be inferred from the trace. However, the service will be less accurate as the reported sanitized location data are far from the actual ones. Tuning Geo-I parameter enables to leverage both privacy protection and service utility.

While the presented methodology can be extended to a large class of protection apps, as stated above, Geo-I will be specifically considered in the remaining of the paper.

3. Control Problem Formulation

The location privacy regulation problem is formulated in this section. The process, actuator, sensors, and disturbance




Figure 3: Protection of a mobility trace using Geo-I. Illustration of a Cabspotting user (abboip) for various configurations of the tuning parameter ϵ .

are defined, as a first step prior to modeling (which will be done in Section 4) and control design (in Section 5). In short, the plant is the protection app, it has two inputs: the user actual location (uncontrollable) and the variance of the noise added for sanitation (the control signal). The plant outputs the sanitized location, on which we build a sensor of utility loss and a sensor of privacy—based on the dispersion of the obfuscated data, to be related to the notion of POI. The disturbance is formally defined as the privacy sensing of the actual location (uncontrollable signal). Justification of the relevance of those signals are given in the respected subsections. A schematic representation of the system is given in Fig. 1. Illustration of the open loop behavior motivating the control problem is further provided.

3.1. System, signals, and sensors

3.1.1. Process

As stated above, the location privacy protection app is considered as the plant. The app takes as input the user actual mobility data and a control parameter (i.e., noise properties), and outputs the sanitized location record. The sanitized data is broadcasted to the Location-Based Service (e.g., navigation or venue finder) and the service response is returned directly to the user. Note that this service loop is out of the scope of the control formulation, and no assumption is made on the impact of the service on the user mobility. Such a system with a measurable output and a tunable input is suitable to be a control plant.

3.1.2. Control Variable

The protection app is assumed to be tunable by at least one signal. In the case of Geo-I, the ϵ parameter can be updated at each iteration and impacts both the POI-oriented privacy and the utility loss, therefore it can be an eligible control signal:

$$u(t) = \epsilon(t) \quad (4)$$

The usual range of values of this control signal is from 10^{-4} m^{-1} to 1 m^{-1} . We warn the reader that the variations of the control signal regarding the obfuscation process is counterintuitive: the smaller the u , the more noise is added (due to the inverse in Eq. (3)). When no noise is applied by the app, we theoretically have u very big. In practice, we set $u = 1 \text{ m}^{-1}$, i.e., a noise of about 1 m that is reasonable given the GPS precision. When $u = 10^{-4} \text{ m}^{-1}$, the noise is of the order of 10 km.

3.1.3. Utility loss sensor

Most location-based service use only the current position to improve users' experience, e.g., weather forecast, venue finders, media tag. The utility loss sensor is then considered as being instantaneous and spatial: the closer the sanitized location sent to the mobile app is to the user's actual one, the better the service will be. The measure of the Utility loss, namely z , is the distance between the sanitized released data 1 and the actual one l [37]:

$$z(t) = \text{dist} [l(t), 1(t)], \quad (5)$$

with $dist[\bullet, \bullet]$ being the Euclidean distance between two points at the surface of the earth. When considering the Geo-I protection mechanism and given Eq. (3), utility loss sensor translates as:

$$z(t) = -\frac{\mathcal{W}(p)}{u(t)}. \quad (6)$$

In this case, the objective of minimizing utility loss is equivalent to maximizing the control signal. The utility objective becomes a constraint on the control signal, we only keep the privacy measure $y(t)$ as a performance signal in the formulation, as illustrated in Fig. 1.

3.1.4. Privacy sensor

With regard to the POI-related privacy notion—to prevent the identification of users' Points Of Interest—we define the privacy sensor based on the spatial dispersion of the locations sent by the user on a past time-window. A small dispersion represents a concentration in space of the obfuscated location records (and also in time due to the time window calculation), which matches with the definition of a POI [26]: the user is perceived as spending significant time in a small area. This signal inversely represents how likely users are to reveal a POI: the higher the privacy measure, the better the protection. Formally, the privacy sensor is defined as being twice the median distance between the location data sent during the time window and the centroid of those points. The centroid l_c of the mobility trace l over a past window of length T is defined as:

$$l_c(t) = \frac{1}{|T|} \sum_{i=t-T}^t l(i). \quad (7)$$

The privacy sensor y of the sanitized location trace l at time t is then:

$$y(t) = 2 \operatorname{median}_{k \in [t-T:t]} (dist[l(k), l_c(t)]). \quad (8)$$

The privacy level is expressed in meters and is to be related with the radius of the smallest POI currently retrievable from the obfuscated trace. The use of median aggregation enhances the privacy signal robustness regarding location measurement noise, such as outlier measures.

3.1.5. Disturbance

Variations in the user mobility (mainly its speed, but also the spatial dispersion of locations, etc.) impact the retrieval of Points Of Interest and thus the privacy sensor y . A user in constant movement, such as in a moving train for example, would have a naturally high privacy, therefore being impossible to extract a POI. Conversely, users wandering around their garden are vulnerable in terms of privacy, and require obfuscation to protect their POIs. These two simple examples illustrate the dependency of the privacy protection level on the user actual mobility data, and more precisely on the *dispersion* of the actual data. We thus use the privacy sensor on the actual location records l to estimate the

disturbance d at time t :

$$d(t) = 2 \operatorname{median}_{k \in [t-T:t]} (dist[l(k), l_c(t)]), \quad (9)$$

with $l_c(t) = \frac{1}{|T|} \sum_{i=t-T}^t l(i)$.

In short, the disturbance is linked to the user actual location data. The disturbance is measurable, has non-uniform sampling (just as the location trace) and presents large fluctuation in terms of both amplitude and frequency (see Section 3.2), making its robust rejection challenging.

3.2. Need for regulation and challenges

Further on, we motivate the need for controlling the protection app parameter to meet a desired privacy level through an open loop behavior analysis. Illustration of the impact of the control signal and the disturbance on the privacy and utility sensors is given in Fig. 4. We use as disturbance a real user mobility trace, the same as in Figs. 2 and 3. Privacy y and utility loss z sensors are measured after applying a constant control signal in two different experiments:

- (i) $u = 10^{-2} \text{ m}^{-1}$ (blue dots), and
- (ii) $u = 10^{-3} \text{ m}^{-1}$ (orange crosses).

Privacy measures are influenced by the user mobility behavior (disturbance). This illustrates the need of a robust dynamic configuration of the protection mechanism. The disturbance d varies across four orders of magnitude and has non-constant sampling: robustness to it is very challenging. Privacy measures show that the control signal u impacts the privacy level. Utility measures z show time variations due to the stochasticity of Geo-I process (i.e., p and θ variables). However, the mean utility loss is constant per trace and only depends on the control signal u . This comforts the choice of defining only the privacy as a reference signal,

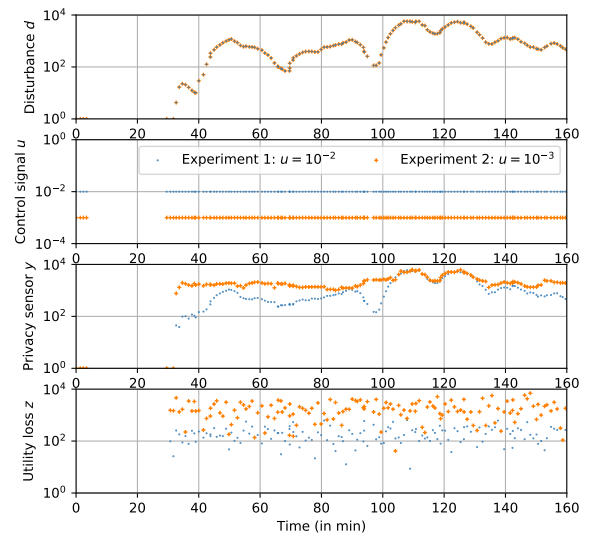


Figure 4: Open loop study: impact of the disturbance and control signal on privacy and utility measured signals. Location record data from user *abboip* of Cabspotting dataset.

and dealing with the utility objective by soundly setting the control variable. Records from 105 min to 130 min illustrate situations where a lower control signal do not increase privacy protection since the user is inherently protected by their movement, it is detrimental for data utility.

It clearly appears here the need for a privacy and utility regulation—achievable by tuning Geo-I’s parameter— that should be robust to the mobility disturbance.

4. Privacy Modeling

Prior to control design, and given the novelty of the problem formulation, extensive modeling efforts are needed. We derive a dynamical model describing and quantifying the app behavior, which links the control signal (noise added) and disturbance (mobility data) to the performance sensor (privacy level). First, we study the static behavior, i.e., without time variation of the inputs. Analytical modeling is performed based on system’s equations, first decoupling the control and disturbance influence, i.e., using the scenario of a stopped user (no disturbance) and then of a moving user without protection (no control action). We further extend the model to tackle both control and disturbance—i.e., privacy protection and movement—through identification from data. The dynamical modeling is eventually performed by identification. The resulting model has a first order dynamics and a non-linear gain. The non-linearities are regarding both the control variable and the disturbance.

4.1. Static Analysis

First, we analyze the static behavior of the system, where the control and disturbance signals are constant. Initially, the respective impacts of the disturbance d and the control signal u on the performance measure, are decoupled. As such, we aim at detecting the presence of non-linearities between u , d , and y .

4.1.1. Decoupling control and disturbance actions

Here we successively study the system without control action and then without disturbance.

Impact of the disturbance on the undriven system Let us consider in a first time the impact of the disturbance on the performance signal, in the case where the control action is null. Note that in our case, this means not modifying the location data before broadcasting it, i.e., from Eq. (3) u tends to infinity.

Theorem 1. *For the static undriven system, one has:*

$$y = d. \quad (10)$$

PROOF. Given Eqs. (8) and (9) and with $1 = l$ (i.e., no control action), one directly has $y = d$. \square

That is to say, in absence of control action, the privacy measure is linear with respect to the disturbance.

Impact of the control signal in an undisturbed system

We now consider the system without disturbance, that is to say the user is stopped, i.e., $d = 0$.

Theorem 2. *For the static undisturbed system with infinitely fast sampling, one has:*

$$y = 10^{a \log(u)+b} \quad (11)$$

with

$$a = -1, \quad b = \log \left(-2 \operatorname{median}_{k \in [t-T:t]} \mathcal{W}(p_k) \right). \quad (12)$$

That is to say, in the absence of disturbance and in ideal recording conditions, the privacy measure is log-linear with respect to the control action. In practice, when the number of records is not large enough in comparison to the selected time window, the offset b varies while the linear gain is always constant $a = -1$.

PROOF. Without disturbance, the user is not moving, we consider without loss of generality that:

$$l(t) = 0, \quad \forall t \in \mathbb{R} \quad (13)$$

The obfuscated locations (Eq. (3)) can thus be written as:

$$1(t) = -\frac{\mathcal{W}(p_t)}{u} \begin{pmatrix} \cos \theta_t \\ \sin \theta_t \end{pmatrix} \quad (14)$$

where p_t is drawn uniformly in $[0, 1)$ and θ_t in $[0, 2\pi)$.

As a first step to the computation of the privacy sensor of Eq. (8), we consider the centroid 1_c of the undisturbed actual location, as:

$$1_c(t) = -\frac{1}{|T|} \frac{1}{u} \sum_{i=t-T}^t \mathcal{W}(p_i) \begin{pmatrix} \cos \theta_i \\ \sin \theta_i \end{pmatrix} \quad (15)$$

where u can be extracted from the sum as in the static case it is independent of the time. Under the assumption that we have sufficiently enough samples, the sum tends to 0 [5]. Indeed, this is expected, as the noise added to the data by Geo-I has a central symmetry. With infinitely fast sampling, one now has:

$$1_c(t) = 0, \quad \forall t \in \mathbb{R}. \quad (16)$$

The distance computation of Eq. (8) is then the norm $|1(k)|$, that given Eq. (14) simplifies as $-\frac{\mathcal{W}(p_k)}{u}$ since by definition $u \geq 0$ and $\mathcal{W}(p_t) \leq 0$. This allows to simplify the privacy sensor of Eq. (8) as:

$$y(t) = \frac{1}{u} \times \left(-2 \operatorname{median}_{k \in [t-T:t]} \mathcal{W}(p_k) \right) \quad (17)$$

Taking the logarithm of the former equation, one has:

$$\log(y) = -\log(u) + \log \left(-2 \operatorname{median}_{k \in [t-T:t]} \mathcal{W}(p_k) \right). \quad (18)$$

\square

An additional challenge of the system comes from the variability of the b parameter. Note that this variability comes from the realization of the random variable p , a behavior designed on purpose in the Geo-I algorithm to harden location trace re-identification. The computation on the time window T with limited samples adds to the variability of b . Asymptotically, we have $b = 0.526$ in average with a standard deviation of ± 0.45 .

4.1.2. Joint impact of control and disturbance

We now take into account the combined impact of disturbance and control action. Given the log-nonlinearity and the offset b highlighted in Theorem 2, we introduce new notations using bold letters.

Notation The updated control signal is defined as:

$$\mathbf{u} = \log(u) - \log(u_L), \quad (19)$$

the performance signal as:

$$\mathbf{y} = \log(y) - \log(y_L), \quad (20)$$

and the disturbance signal as:

$$\mathbf{d} = \log(d) - \log(y_L). \quad (21)$$

where u_L is a working point (typically the center of the control signal range of variation) and y_L its corresponding privacy level (using results of Theorem 2):

$$\log(y_L) = a \log(u_L) + b \quad (22)$$

Note that d is a measure using the privacy sensor (see Section 3.1.5) and is thus linearized using y_L . The use of the logarithmic function in Eq. (19) ensures that the control signal \mathbf{u} can take values in \mathbb{R} : the constraint of non-negativity on u has been removed.

Asymptotic behavior From Theorems 1 and 2, one can retrieve the asymptotic behavior of the performance signal \mathbf{y} depending on the values of \mathbf{u} .

Lemma 3. *The asymptotic equations of the privacy \mathbf{y} in presence of a control action \mathbf{u} and a disturbance \mathbf{d} are:*

$$\lim_{\mathbf{u} \rightarrow -\infty} \mathbf{y} = \mathbf{a}\mathbf{u}, \quad \text{and} \quad \lim_{\mathbf{u} \rightarrow +\infty} \mathbf{y} = \mathbf{d}, \quad (23)$$

with the bound value, at which \mathbf{y} switches of asymptotic behavior, being: $\mathbf{u}_0 = \frac{\mathbf{d}}{\mathbf{a}}$.

That is to say, the output \mathbf{y} is linear with respect to the control input \mathbf{u} for small values of \mathbf{u} , and constant at a disturbance-dependent level for large values of \mathbf{u} . The bound depends on the disturbance level.

PROOF. When $\mathbf{u} \rightarrow -\infty$, one has from Eq. (19): $u \rightarrow 0$, and so $\frac{1}{u} \rightarrow +\infty$. Thus, using Eq. (3), l become negligible and one has:

$$\lim_{\mathbf{u} \rightarrow -\infty} \mathbf{1} = -\frac{\mathcal{W}(p)}{u} \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix},$$

that, following Eq. (14), corresponds to an undisturbed system. Then, using Theorem 2, one has:

$$\log(y) = a \log(u) + b \quad (24)$$

Using the notations of Eqs. (20) and (22), this translates as:

$$\begin{aligned} \mathbf{y} &= a \log(u) + b - \log(y_L) \\ &= a \log(u) + b - a \log(u_L) - b \\ &= \mathbf{a} \mathbf{u} \end{aligned} \quad (25)$$

Conversely, when $\mathbf{u} \rightarrow +\infty$, that is to say for $\frac{1}{u} \rightarrow 0$, one has from Eq. (3):

$$\lim_{\mathbf{u} \rightarrow +\infty} \mathbf{1} = l,$$

that corresponds to an uncontrolled system. Thus, using Theorem 1, one has $y = d$. Given Eqs. (20) and (21), this rewrites as $\mathbf{y} = \mathbf{d}$. The bound on \mathbf{u} can be found using the value \mathbf{u}_0 at which the two asymptotic equations meet:

$$\mathbf{y} = \mathbf{a} \mathbf{u}_0 = \mathbf{d}. \quad (26)$$

□

Transition between asymptotic trends To practically capture the complex transition behavior of the static characteristic, we opt out for identification from data.

We experimentally retrieve the static characteristic. Several experiments are conducted with constant control signal and disturbance, the constant values being different between experiments. Average privacy is computed for each experiment. The control signal values u are chosen as log-uniformly distributed in its definition range. A constant disturbance is characterized by a constant dispersion of actual data (see Section 3.1.5), expressed by a constant speed in our experiments. The speeds are taken with various values: high, low or null (the user is stopped).

An illustration of a static characterization is given in Fig. 5. Those results illustrate Theorems 1 to 3: (i) the logarithm of the privacy measure is linear with respect to the logarithm of the control signal for low values of u and (ii) for large values of u there is a saturation at a constant value, whose level depends on the disturbance.

The smooth transition between linear and saturated parts can be identified from Fig. 5 as deriving from the following equation:

$$y = \frac{u_d}{u} d \sqrt{1 + \left(\frac{u}{u_d}\right)^2}, \quad (27)$$

with

$$u_d = 10 \frac{\log(d) - b}{a}. \quad (28)$$

Analytical asymptotic results of Theorem 3 can be retrieved when varying the relative value of u compared to u_d (i.e., compared to d). Experimental validation is given in Section 6.1.

Finally, we retrieve the static gain from Eq. (27).

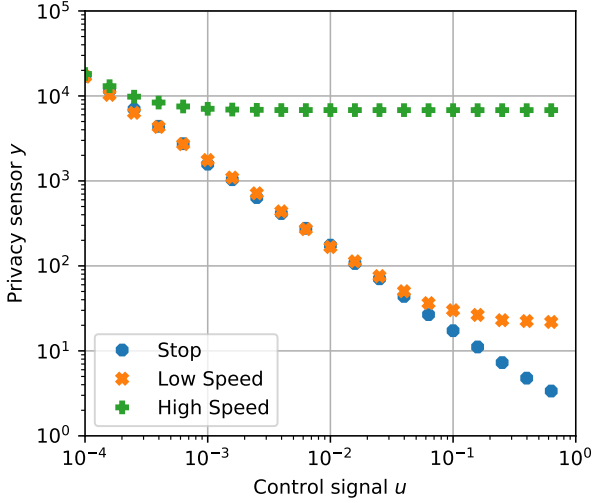


Figure 5: Static characteristics of the control input to privacy output transfer function for various (constant) disturbance scenario. Each data point is the averaging over a whole experiment. Data extracted from Privamov user 51.

Theorem 4. *The gain K between the controlled obfuscation \mathbf{u} and the privacy level \mathbf{y} is non-linear, and depends on both the control \mathbf{u} and the mobility disturbance \mathbf{d} :*

$$K(\mathbf{u}, \mathbf{d}) \triangleq \frac{d\mathbf{y}}{d\mathbf{u}} = \frac{-1}{1 + 10^{2\left(\mathbf{u} - \frac{\mathbf{d}}{a}\right)}} \quad (29)$$

For readability considerations, proof of Theorem 4 is in Appendix A. The gain is an input non-linearity in the sense of the Hammerstein-Wiener modeling [52]. It is asymptotically consistent with previous results.

4.2. Dynamic modeling

The impact of control signal time variations on the system is now studied. Black-box identification is performed on a stopped user, i.e., without disturbance, to ensure being in the linear zone of the static characteristic. The dynamics of the system are expected to come only from the time window calculation of the privacy signal, which motivates the time analysis over the frequency one. An input step variation is applied in which the values of the initial and final level are chosen in the linear range of the system ($u < u_d$), as in the constant part the control signal has no impact on the privacy level. The dynamic model captured with such methodology gathers the plant (protection app) and the privacy sensor processes.

The evolution of privacy through time is reported in Fig. 6. The relation between the control input and the privacy measure can be approximated as a Hammerstein-Wiener model, with a first order LTI transfer function and an input non-linearity—derived from Eq. (29):

$$\mathbf{H}(s) = \frac{\mathbf{Y}(s)}{\mathbf{U}(s)} = \frac{K(\mathbf{u}, \mathbf{d})}{1 + \tau s} \quad (30)$$




Figure 6: Step response of the undisturbed system in its linear zone. Privamov user 51.

s being the Laplace variable and with $\tau = \frac{T}{3}$, as the rise time corresponds to the length of the privacy metric time window T . The non-linear gain $K(\mathbf{u}, \mathbf{d})$ reflects the saturation effect at large control-signal values.

5. Control Strategy

The Location Privacy controller is presented, starting with the formulation of the control objectives and followed by a detailed description of each control component.

5.1. Specifications

The controller aims at tackling the challenge of personalization of protection mechanisms through dynamical adaptation to users' mobility practically. In details, users' first objective is to keep acceptable privacy levels despite their highly varying mobility patterns. In control terms, this objective translates in rejecting the users' movements disturbance. Users can additionally specify a desired privacy protection level (a minimal threshold), corresponding to a reference tracking objective. This reference can vary significantly, depending on if the user is in a utility-sensitive or privacy-sensitive situation. The controller is desired with zero steady state error, small overshoot (e.g. less than 10% [39]), and a reasonable settling time with respect to the user mobility (i.e., of T). Indeed, the control does not aim at increasing the rapidity of the system, as it would result in a non-realistic mobility trace. The controller is aimed to be implemented on a smartphone or any other mobile device. It thus needs to show limited overhead in terms of computation, storage, and communication to have a limited impact on users' experience and device battery. Given the very large audience susceptible of being interested in privacy




Figure 7: Control loop schema showing the three control components: a feedback, a pre-compensation and anti-windup

protection mechanisms, user-friendliness is a key feature: very few inputs from the user should be asked. Eventually, as the location is shared, first and foremost, for the user to benefit from a service, the obfuscation should not degrade significantly data utility.

To sum up, the closed-loop specifications are:

- (i) reject the mobility disturbance,
- (ii) follow the privacy reference and increase as much as possible service utility when privacy constraints are met,
- (iii) induce low overhead and ask for limited users inputs.

The controller is thus composed of a feedback controller (see Section 5.2) for reference tracking and to react to the presence of disturbance, a pre-compensation action (detailed in Section 5.3) to anticipate reference sharp variations, and an anti-windup mechanism (see Section 5.4) to prevent irrationally losing service utility. Low-complexity versions of those controllers components are chosen to meet the third objective, and an emphasis is made on finding universal control parameters to avoid any tuning by users. The control loop is illustrated in Fig. 7, showing the articulation of the three control components. As motivated in Section 4.1.2, the signals are in their linearized form. We define y_{sp} as the linearization of the desired privacy level y_{sp} :

$$y_{sp} = \log(y_{sp}) - \log(y_L). \quad (31)$$

The user objective is expressed in meters, i.e., a value of $y_{sp} = 100$ m means that the user does not want POIs of 100 m diameter or smaller to be retrievable.

5.2. Feedback Controller

A feedback action is used for reference tracking. Disturbance rejection is also ensured using a feedback controller, as it reacts to the impact of the mobility disturbance of the user on the privacy values, and compute accordingly a compensatory control signal. A PI controller is used for zero steady state error, thanks to the integral action, while being sufficiently robust to the stochasticity of the plant (i.e., the app) and to the effects of the disturbance.

The PI part of the controller is expressed as:

$$PI(s) = \frac{U^{PI}(s)}{Y_{sp}(s) - Y(s)} = \frac{K_I}{s} + K_P. \quad (32)$$

Tuning guidelines. The parameters are tuned using pole placement as detailed in [6]:

$$K_I = \frac{\tau}{K_L \cdot \tau_{obj}}, \quad K_P = \frac{1}{K_L \cdot \tau_{obj}}, \quad (33)$$

with τ_{obj} the pole of the objective closed loop, fixed by the desired response time. In our privacy use-case, the user sets its desired POI duration time T in the definition of the privacy signal (see Eq. (8)). To have a control loop faster than the system, we set $\tau_{obj} = \frac{\tau}{3}$. The linear gain is used: $K_L = a$.

The PI controller is discretized to cope with the non-constant sampling time:

$$u^{PI}(t_i) = (K_I(t_i - t_{i-1}) + K_P) e(t_i) - K_P e(t_{i-1}) + u^{PI}(t_{i-1}) \quad (34)$$

with $e(t_i) = y_{sp} - y$ being the error.

5.3. Pre-compensation Control

A pre-compensation action is added to enhance reference tracking. The reference signal can present large and sudden variations. This happens for instance when the user urgently needs a service (such as starting navigation) or conversely finished using it (end of the course). The pre-compensation control anticipates the effect of such sudden reference changes. Using a purely reactive control would cause significant drifts during the sharp reference variations.

We use a pre-compensation action as an anticipation control, able to deal with the changes of the reference signal:

$$u^{PC}(t+1) = \frac{y_{sp}(t)}{K} \quad (35)$$

Both pre-compensation and feedback controllers are thus soundly tuned without the need for users inputs.

5.4. Anti-windup

The anti-windup action has two purposes: *i*) ensuring that the control signal will not take extremely low values which would result in tremendous data distortion, and *ii*) ensuring that it does not reach significantly high values for which the utility gain is negligible whereas the controller would need too much time to react and decrease the control

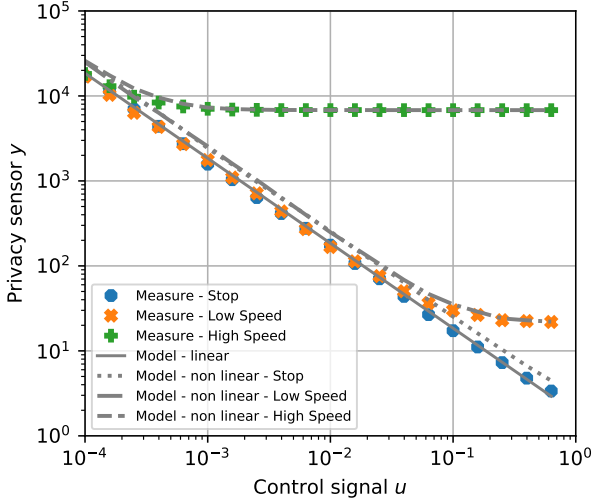


Figure 8: Static modeling evaluation: the non-linear static gain allows for a fine fit to experimental data. Privamov user 51.

signal. To ensure such behavior, an anti wind-up strategy is added as well as an actuator saturation [49]:

$$\begin{cases} \mathbf{u}^{\text{PI}} = \min(\max(\mathbf{u}^{\text{PI}}, 2\bar{\mathbf{u}}), 2\mathbf{u}), \\ \mathbf{u} = \min(\max(\mathbf{u}^{\text{PC}} + \mathbf{u}^{\text{PI}}, \bar{\mathbf{u}}), \mathbf{u}). \end{cases} \quad (36)$$

where thresholds $\bar{\mathbf{u}}$ and \mathbf{u} are chosen according to Geo-I's common range of variations, ensuring the stability of the system. Note that the PI action has a wider range of variation allowed, as the pre-compensation may significantly shift the global control action.

6. Evaluation

We first present the validation of the model and controller using real-life data. Evaluation of the generality of the control to several users from various datasets is then performed. Analysis of utility preservation and overhead ends the section with practicality concerns.

Models and control parameters values are summarized in Table 1. We consider that the time window (T in the privacy definition) is to be set by the users so that their points of interest are hidden. In our case, we chose to protect places where the users stay longer than $T = 15$ min. This avoids insignificant stops such as waiting at a traffic light to be obfuscated, while being ambitious on the points of interest to

Parameter	Value	Parameter	Value
T	15 min	u_L	10^{-2} m^{-1}
a	-1	y_L	335.7 m
b	0.526	τ_{obj}	15 min
K_L	-1	K_P	-0.334
τ	5 min	K_I	-0.0011

Table 1: Values of all the parameters for the model and control algorithm




Figure 9: Dynamic model validation: robustness to user's mobility. Experimental data in blue and proposed model in orange. Privamov user 51.

protect: not only work and home places but also information such as shopping, meeting, or worship places. All others parameters are drawn from theory developed in Sections 4 and 5, linearization of the control signal is done at the center of its definition range.

Experiments are carried out using Python 3 on a laptop equipped with an Intel Core i7-1185G7 CPU clocked at 3GHz x 8, running under Ubuntu 20.04 LTS. Codes for reproducing experiments and analysis are made open [11].

6.1. Model validation

We evaluate the accuracy of the model in capturing the privacy level of a user through time, knowing the control input (protection app parametrization) and the disturbance (user's movement). We consider the non-linear model of Eq. (29), and compare it with the linear asymptotic gain $\mathbf{K} = a$ as in Theorem 3. First modeling performance is shown without considering time influence (static scenario), then the prediction accuracy in real-time is presented.

6.1.1. Static privacy prediction

The static characteristic shows the impact of the control signal u on the privacy value y . Experimental results of Fig. 5 are compared to the models predictions in Fig. 8. The parameter b is found using linear regression (implemented through sklearn in Python): $b = 0.26$. The linear model (continuous line) successfully captures the system's behavior for small values of the control signal. However, when the user is moving at high speed, the linear model performs correctly on less than a fourth of the control input range. The non-linear model (dashed, dotted and dashed-dotted lines depending on the value of the disturbance) successfully captures the saturation effect. The smooth transition between linear and




Figure 10: Dynamic model evaluation: handling control signal variations. Experimental data in blue and model in orange. Privamov user 51.

saturated zone is correctly modeled, allowing for a refined control action around the non-linearity breakpoint. Results of Fig. 8 favor the use of the non-linear model over the linear one, as in the following dynamic modeling evaluation.

6.1.2. Dynamic privacy prediction

This section investigates the dynamic aspect of the model, i.e., its ability to predict in real-time the privacy level of a user.

Fig. 9 presents the privacy measured (blue dots) and modeled (orange line) through time in the top plot; for a constant control signal (set at $u = u_L$, middle plot) and under a challenging disturbance (bottom plot). A 24-hour mobility trace of the user 51 from Privamov dataset (chosen randomly) is used as experimental data. The disturbance presents both high and low frequency variations, as well as non-constant sampling time, as can be seen around 9100 min. The predicted privacy value (continuous orange line) corresponds to the measured privacy (blue markers) with very good accuracy, despite the highly varying disturbance. Dealing with non-constant sampling is however difficult, as highlighted by the occurrence of small oscillations of the model prediction around 9100 min.

In a second time, the control signal is randomly varied, and the model performances are illustrated in Fig. 10. The random control signal is generated with various amplitudes ($u \in [-2, 2]$) and periods (from 10 s to 30 min), and is reported on the middle plot. The top plot presents the privacy levels: one can see again a precise fit between the measured values (blue crosses) and the prediction (orange line). We see again here the impact of non-constant sampling through the

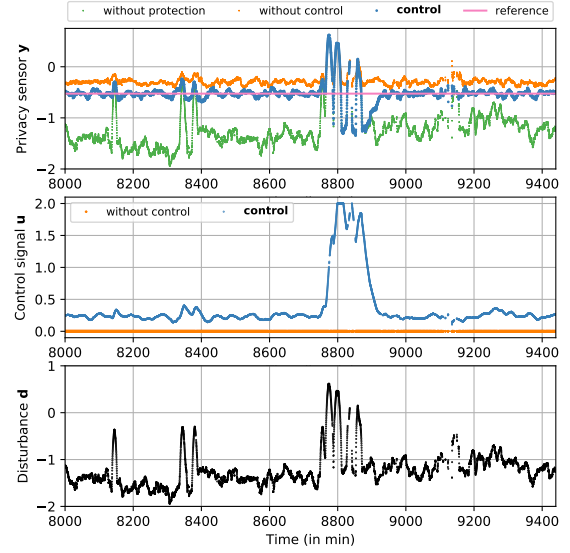


Figure 11: Controller validation: robustness to user's mobility. Comparison of our proposed control strategy with a static protection app configuration (i.e. without control) and a non-protected user (i.e. without protection). Note that the control signal plot shows a constant signal at $u = 0$, corresponding to the constant parametrization of Geo-I **without control**. Privamov user 51.

appearance of high frequency behavior. However, these inaccuracies may not be detrimental to the control performance, which supports the use of the non-linear model.

To sum-up, the modeling is fairly accurate both in steady and in dynamical state.

6.2. Analyzing Controller's performance

The controller performance is evaluated with regard to the objectives presented in Section 5. First the robustness to the user's mobility disturbance is studied, then its ability to follow a privacy reference specification is evaluated. Utility and control overhead are discussed respectively in Sections 6.4 and 6.5.

In all the controller evaluation experiments, results of several strategies are presented for comparison. The privacy level achieved with our **control** strategy is compared to the privacy level of the user **without protection**. We recall that this level is not null, as the user's speed inherently provides some protection. Comparison is also given with the privacy achieved when using a protection app with a static configuration at $u = 0$, i.e., **without control**. Note that this strategy corresponds to the state-of-the-art protection mechanism Geo-I. The comparison with the scenario **without protection** enables to show moments when no obfuscation is needed, while the comparison with the scenario **without control** illustrates the benefits of a dynamic configuration of the protection app.




Figure 12: Control validation: handling privacy with reference variations. Note that the control signal plot shows a constant signal at $u = 0$, corresponding to the constant parametrization of Geo-I **without control**. Privamov user 51.

6.2.1. Robustness to disturbances

First, the controller's robustness to the presence of a highly varying disturbance is evaluated. The privacy reference value is fixed and set at $y_{sp} = 10^2$ m, i.e., $y_{sp} = -0.526$. Results are given in Fig. 11: the top plot presents the privacy signals of the different scenarios (without protection, without control and with our control) and the constant reference; the middle plot shows the control signal; and the bottom plot the disturbance.

During most of the experiment, the user has a privacy level **without protection** lower than the reference. In those moments, the **control** leverages the control signal to meet the reference with a good precision. **Without control**, the privacy level is high and close to the reference, however with a static error.

Around 8000 min, the disturbance oscillates with large amplitudes: the system **without protection** shows an increased privacy. The **control** signal is then increased by the controller up to the upper bound. Remind that a high control value means little noise added to the mobility data (due to the inverse formulation in Eq. (3)), which results in utility savings. The scenario **without control** but with a constant u enables to preserve a high privacy level throughout the experiment, at the cost of constant data obfuscation, which has a high utility cost (see Section 6.4 for quantified details)

Those results validate the use of the controller to be robust to users' mobility with a limited utility cost.

6.2.2. Reference tracking

Second, we evaluate here the ability of the controller to achieve user's time-varying reference privacy level.




Figure 13: Control validation on different users. Privamov 14 user. Fast and uniform sampling provides smoother control.

Fig. 12 shows experiments with the same disturbance (i.e. we use the same mobility trace previously presented) whereas with a varying reference scenario in this case (pink continuous line). The reference consists of a random signal generated with varying amplitude ($y_{sp} \in [-2, 2]$) and period (from 5 min to 5 h).

When the privacy reference value is higher than the privacy **without protection**, the controller increases the control signal so that the privacy with **control** meets the reference value. The reference tracking is (i) *stable*, despite the modeling high frequency variations; (ii) *precise*, the static error is null even if some noise is visible due to the system stochasticity; (iii) *fast*, the rise time is of the order of 10 min while the settling time can reach 50 min, and (iv) with some *overshoot*, which amplitude varies from about 10% at 9150 min to 40% at 8150 min. Rapidity and overshoot could be improved, for example using more advanced control techniques. The scenario **without control** is not able to follow the reference when it is too high (see for instance from 8000 min to 8580 min) while uselessly degrading the service quality at other periods (see for instance 8920 min to 9170 min).

Overall, the presented controller manages to precisely follow the dynamic reference and reject the mobility disturbance with a user-independent configuration, a low computation overhead and limited control cost.

6.3. Robustness: evaluation on different users and datasets

The controller's ability to adapt to other users' behavior without any change on the control algorithm (i.e., with the same tuning) is now evaluated.

A similar approach of random reference signal is applied to another user of the same dataset (Fig. 13) and two users




Figure 14: Control validation on different users. Cabspotting abboip user.

from another dataset (Figs. 14 and 15). As Section 6.2 showed the advantage of our **control** compared to the scenario without control, for improving the reading, we show only the comparison to the scenario **without protection**—as it enables to explain why some low reference values cannot be tracked. Globally, in all experiments the **control** privacy reached the desired reference value while **without protection** the privacy protection level is poor.

The user of Fig. 13 shows a significantly different scenario, with a strong disturbance (low privacy **without protection**) due to its low speed (in average around 5 km h^{-1} , i.e., someone walking). The **control** shows high performance in terms of precision in this case, and little static noise. This user illustrate a second challenge: dealing with record absence, i.e., anomaly large sampling period, see 945 min to 975 min. Despite such extreme conditions, the **control** stays stable and converges to the reference in a few tens of minutes.

We now evaluate our control with the data of users from another dataset, i.e., with different recording conditions. The disturbance signal from Fig. 14 illustrates those differences: *i*) the sampling time is larger (30 s in average) and *ii*) the users move faster (as they are in cars), leading to higher values of disturbance. Despite the large sampling periods, i.e., fewer control actions, the **control** keeps the privacy level at the reference value with good precision. The control rapidity is similar, while the steady state presents large oscillations due to large oscillations in the disturbance signal. The privacy level **without protection** is globally high, with moments where it is higher than the reference, see for instance 420 min to 600 min. In those cases, the **control** sets the input to its maximum value, meaning very little noise (1 m) is added to the data, so utility is preserved. Note that




Figure 15: Control validation on different users. Cabspotting oilrag user.

this behavior is not detrimental to our use-case, as utility is maximal and privacy is high (even above the reference).

The last user of our comparison, in Fig. 15, also presents long sampling periods and high disturbance values; while additionally *i*) data are not collected for a long period, from 80 min to 190 min, and *ii*) the disturbance shows very large and quick variations. Even in those extreme conditions of few control actions possible and extremely long sampling period, the **control** stays stable and precisely tracks the reference. The large and rapid variations of the disturbance and privacy **without protection** are challenging for our controller, that however manages to keep tracking the reference while with significant noise around the desired value.

Those results show the generality and robustness of the controller that is able to perform correctly, with the same parameters, with various users from different datasets. Our controller has shown to be practical thanks to its universality, while being able to adapt to time variations, thus allowing for personalized privacy protection. In addition to the user behavior (disturbance variation profile), the recording conditions (non-constant sampling with large non recording periods) also affect the controller. Despite those challenges, our control reaches fast and accurate performances in favorable conditions (stable and fast sampling rate) while being stable and still performing fairly good in degraded conditions.

6.4. Utility preservation

After validating the privacy-related performance of the controller, the preservation of data utility is evaluated. Beforehand, the methodology and metrics are presented.

User	Configuration strategy	Distortion z (in m)		Service Utility $z_{service}$ (in %)			
		median	90 th percentile	NavigationLyon (thd=10 m)	NavigationSF (thd=100 m)	VenueFinder (thd=1 km)	WeatherForecast (thd=10 km)
51	control	167.3	404.3	1.3	27.5	99.9	100
	without control	165.5	390.2	0.5	27.0	99.9	100
14	control	172.6	414.5	1.8	25.3	99.7	100
	without control	169.6	393.0	0.4	26.2	100	100
abboip	control	2.7	162.1	71.6	85.1	98.9	99.3
	without control	166.1	375.7	0.2	24.9	99.3	99.3
oilrag	control	2.0	47.7	83.1	90.0	97.4	97.6
	without control	163.8	378.5	0.0	24.7	97.7	97.7

the lower the better
 the higher the better

Table 2: Utility Evaluation for all users. Best results are in bold.

6.4.1. Metrics

Several indicators are used to capture the diversity of the utility notion, both theoretically and in practice. First, we investigate the spatial distortion of data, as a general but theoretical utility notion. Then, we compute service-based utility metrics, that evaluate the functionality of different location-based services (such as navigation or weather forecast) based on our sanitized data. Details of their computation are given hereafter.

Firstly, we evaluate data distortion, that allows comparison with state of the art. More specifically, we compute the instantaneous spatial distortion between the original location and the obfuscated one, z , following equation (5). The median value and the 90th percentile (worst-cases) aggregated over the whole trace are computed [37] to give a per-user utility evaluation. While those metrics enable to have an overall idea of the data distortion, they are not able to capture *service* quality loss [2]. Most location-based services present a threshold-based behavior: if the distortion is lower than the service-specific threshold, the service can still work, while if it is higher, the service gives useless results. As an example, in a dense European city such as Lyon (place of Privamov records), a Navigation app can provide the desired route if data spatial distortion is lower than about 10 m; while the app may give a wrong route if distortion is higher, resulting in zero utility. A Weather Forecast service presents the same behavior: the forecast is accurate for data with a distortion below around 10 km, while above this threshold, the weather can considerably change. To capture this binary threshold-based behavior, Service Utility metrics $z_{service}$ are defined for 4 representative services. $z_{service}$ gives the proportion of the time a given service worked correctly, computed based on a threshold thd on data spatial distortion:

$$z_{service} = \text{mean } \mathbf{1}_{X < thd}(z) \quad (37)$$

In addition to the above-mentioned examples, a Navigation app used in San Francisco is considered (record place of Cabspotting), a city in which blocks are of much larger size than in Lyon, thus a threshold value of 100 m. A Venue Finder app completes the metrics, for which a spatial distortion threshold of 1 km is used.

6.4.2. Evaluation

Distortion and service-based metrics are computed on the data from the experimental evaluation presented in Sections 6.2 and 6.3. Results are given in Section 6.4.2. For each user, we compare the scenario with control (our approach) and without control (static configuration with $\mathbf{u} = 0$, i.e., state of the art). For a fair comparison, we set $\mathbf{y}_{sp} = 0$ so that to have similar privacy performance in both scenarios.

First, we analyze the spatial distortion, for which lower values are better. Significantly lower values are achieved with the control than without, for all users. For *abboip* and *oilrag* users, the very low values of the medians of about 2 m show negligible control impact half of the time. Indeed, when the user is moving, no protection is needed, so no noise is added to the data; which guarantees a perfectly efficient service. Without control, a constant noise is nevertheless applied, resulting in significant distortion and loss of utility. Users 14 and 51 show higher distortion, due to their profile with a strong disturbance (i.e., low speed, easily extricable points of interest). Note that for 90% of the data, the distortion with control is significantly lower than without for *abboip* and *oilrag* users. For the other users, results with and without control are comparable: less 400 m distortion for 90% of the data points.

Service utility metrics complete the analysis with concrete impact on location-based services. Low demanding services such as weather forecasting and venue finder are almost always usable ($> 97\%$ of the time), for all users and both with and without control. The most demanding service, i.e., navigation in Lyon, is functional more than 70% of the time for the users with low disturbance (*abboip* and *oilrag*), while the service cannot be used at all for the strongly disturbed users (the service works less than 2% of the time). On those demanding services, the controller ensures significantly better utility than without.

As a conclusion, the unavoidable utility degradation inherent to the privacy protection can be observed, while during most of the trace and for most services, the utility is preserved with our control approach. In comparison to a static protection without control, the service quality levels achieved with the dynamic control are significantly higher, especially for users with a high mobility. Our controller is

practically usable, without much service degradation, while preserving privacy.

6.5. Control overhead

The overhead due to the use of our control is evaluated according to three aspects: computation, storage and communication. We first evaluate the overhead of *control*, i.e., the dynamic decision on parametrization, and then the one of the obfuscation algorithm—even if it is state of the art.

The computation complexity of the controller, defined by Eqs. (34) and (35), is $O(1)$. Experimental tests have shown a computing time of the magnitude of the millisecond for a control decision, that is negligible in regards with the sampling period of the datasets (10 s to 60 s). The impact on storage is also negligible, as our first order controller only requires to store data from the last sampling period. Eventually, the impact is null on communication, as only one data point is sent to the service, exactly as in the scenario without control. The computational complexity of Geo-I is higher than our control, as it consists in drawing two random variables and apply the Lambert W function, while it stays largely negligible. Geo-I has no communication or storage overhead.

To conclude, the controller introduces very few computation and storage overhead and no extra communication. This approach is therefore suitable for a practical implementation on a smartphone, with guarantees of fast execution and low battery usage. Other control algorithms with higher computational complexity could thus be used, they will be investigated in future work.

6.6. Discussion

We now discuss the performance and limitations of our proposed protection mechanism in front of a set of different attacks and analytics. We first consider the case of a reidentification attack based on POIs [43]. Our approach completely obfuscates POIs which diameter is smaller or equal to the reference value. Knowing the attack parameters, one can thus defend against it by soundly setting the reference signal. Let us now consider an attack aiming at reconstructing transportation means. It is based on analytics on the speed and acceleration of the users. With our approach, random noise is added to the positions, which means that speed and accelerations are distorted. No valuable information can thus be extracted on the transportation means. A popular analytics of mobility data consist in building heat maps [36]. Note that those heat maps can be seen both as privacy attacks (to perform reidentification for instance) and service utility (e.g. for traffic analysis). Depending on the relative value of the heat map granularity and on the privacy reference set in our framework, this analytics can be distorted or preserved. Eventually, let us consider the case of filtering of the mobility data, e.g. using a Kalman filter. Our approach adds dynamically varying noise on data, which can be filtered out by a well-performing attack.

As for future works, other control approaches could be investigated, allowing for instance to preserve some utility-related features of mobility data (e.g. speed for a navigation

app). Optimal approaches would enable to remove the need for stochasticity in the noise, thus providing robustness to filtering attacks.

7. Related Work

We situate the contributions of our work with respect to two research areas: on deriving privacy metrics computable online, and on adapting protection mechanisms in time, space, and for each user.

7.1. Online privacy measures

Privacy metrics are of three main types: [43] (i) extracted from formal guarantees (like k -anonymity or differential privacy), (ii) computed from attacks, or (iii) based on data-distortion. Formal guarantees require knowledge of a dataset with different users, which make those metrics unusable in our user-based, online scenario. Attack-based metrics evaluate the accuracy, correctness and certainty of a given privacy attack on data [45]. They present three main limitations: being specific to a particular attack, high complexity thus low practicality, and unfeasible privacy protection with decent utility. In particular, efforts have been made toward online metrics by considering the time correlation between data [53], however at the cost of exponential complexity [54]. While bounding techniques allows for more efficient computation, the complexity of this approach is much higher compared to our POI-based metric. Defeating a well-performing privacy attack has been shown infeasible in practice if users want to keep a decent data utility [32]. Alternative metrics are needed that can capture the utility to privacy trade-off.

We rely on a privacy metric based on data dispersion, i.e., POI, that belong to the data distortion metrics. State-of-the-art POI metrics are offline [22, 19, 33, 9], they require the whole trace to be computed. Our paper contributes in presenting the first privacy metric based on POI that can be computed online. Our approach also has the benefit of being general, and not optimized for a specific attack foiling.

7.2. Dynamic location protection

The location privacy literature proposes few works on the protection apps configuration challenge, most of them treating the problem as a static one—i.e., working on already collected databases, possibly of temporally uncorrelated locations—and able either to work on a specific mechanism [14, 3] or to choose between several [42, 10]. L2P2 [51] is a dynamic objective-driven protection configuration law for location privacy. It leverages the size of the cloaking area in which the location is reported. Even if the configuration adapts to the changes in the reference privacy, the algorithm does not take into account the user's movements, unlike our approach. The control literature regarding differential privacy is mainly focused on designing attacks using state or input observations [17, 29]. Few works have addressed the configuration challenge of the differential privacy parameter [16, 12, 31], but always in a static way, i.e., not taking time variations into account. A feedback control approach

for privacy has been introduced by Wang et al. [50], considering a general differential-privacy mechanism and an attacker performing state observation. However, this work is not specific to the location scenario and the controller (a unit gain feedback) is not considered from a robustness perspective, as the mobility application requires. A dynamic sampling method to protect the release of vehicle's real-time trajectory data has been proposed [25], with drawback regarding temporal distortion, imposing to the user a delay of the service. Personalized privacy protection can be achieved using a user profile (assuming *a priori* knowledge of the user's POIs) [38]. This user profile is used to adapt geographically GeoI's noise *distribution*, i.e., not its parametrization but changing the Laplacian distribution for an optimal one. The main limits of this work are the assumption of *a priori* user knowledge (detrimental to usability in practice), time is considered as a series of service queries (and not as a regular location broadcast), and results are restricted to a 1D scenario. Adaptive location preserving privacy mechanisms [4, 28] adjust the amount of noise required to obfuscate the user's location based on the correlation level with its previous obfuscated locations. Those works take as privacy metric the predictability of one's mobility, which is orthogonal and complementary to our POI-based approach.

8. Conclusion

This work tackles the challenge of robust dynamic privacy protection of a mobile apps' user. This online scenario is particularly sensitive to privacy attacks if a malicious agent has access to the real-time position of the user. Focus is made on the protection of user's points of interest, an indicator of behavior and identity. Protection mechanisms from the literature come with an unavoidable reduction of the service utility, as information is deteriorated to ensure privacy. We tackle three challenges of such protection apps: (i) their usability in practice by non experts, especially regarding their configuration, (ii) the possibility to dynamically change one's privacy level requirements through time, and (iii) the robustness to users' mobility specificities to have personalized protection. We present a control-based approach enabling users to control their privacy when using such protection mechanisms, while keeping an eye on utility loss, regardless of their mobility patterns. Contributions are on the novel problem formulation and particularly a definition of real-time (points of interest oriented) privacy metric; on the non-linear modeling of the system; and on a control strategy with universal configuration. Evaluation carried out using real data highlights the performance and robustness of the controller for all users, with high service utility preservation and low computational overhead. Control algorithm and evaluation codes are available online [11]. Further directions of study can suggest the presentation of the control problem as a global optimization problem solvable online or where prediction of the mobility patterns could be taken into account, e.g., using machine learning.

Acknowledgment

The authors warmly thank Mirko Fiacchini (Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, France) for the enriching discussions and revision of the paper.

A. Proof of Theorem 4

PROOF. From Eq. (20) and given Eq. (27), we have:

$$\mathbf{y} = \log \left(\frac{u_d}{u} d \sqrt{1 + \left(\frac{u}{u_d} \right)^2} \right) - \log y_L \quad (38)$$

Simplifying the logarithm of products and powers, one has:

$$\mathbf{y} = \log(u_d d) - \log u + \frac{1}{2} \log \left(1 + \frac{u^2}{u_d^2} \right) - \log y_L \quad (39)$$

Using the identity $u = 10^{\log u}$ and expressing \mathbf{u} from u (Eq. (19)):

$$\begin{aligned} \mathbf{y} &= \log(u_d d) - \mathbf{u} - \log u_L \\ &+ \frac{1}{2} \log \left(1 + \frac{1}{u_d^2} 10^{2\mathbf{u}+2 \log u_L} \right) - \log y_L \end{aligned} \quad (40)$$

Now we compute the non-linear gain by deriving \mathbf{y} according to \mathbf{u} . Several constant terms (relative to \mathbf{u}) cancel out:

$$\begin{aligned} \mathbf{K} &\triangleq \frac{d\mathbf{y}}{d\mathbf{u}} \\ &= -1 + \left[\frac{1}{2} \log \left(1 + \frac{1}{u_d^2} 10^{2\mathbf{u}+2 \log u_L} \right) \right]' \end{aligned} \quad (41)$$

to derive the function composition, we use a temporary notation $f(x) = \log \left(1 + \frac{x}{u_d^2} \right)$ and $g(x) = 10^{2x+2 \log u_L}$. Given that $f'(x) = \frac{1}{\ln(10) \times (u_d^2 + x)}$ and $g'(x) = 2 \ln(10) \times 10^{2x+2 \log u_L}$:

$$\begin{aligned} \mathbf{K} &= -1 + \frac{1}{2} f'(g(\mathbf{u})) g'(\mathbf{u}) \\ &= -1 + \frac{2 \ln(10) \times 10^{2\mathbf{u}+2 \log u_L}}{2 \ln(10) \times (u_d^2 + 10^{2\mathbf{u}+2 \log u_L})} \end{aligned} \quad (42)$$

After simplification of $2 \ln(10)$ terms and grouping terms under the same denominator, one has:

$$\begin{aligned} \mathbf{K} &= \frac{-u_d^2}{u_d^2 + 10^{2\mathbf{u}+2 \log u_L}} \\ &= \frac{-1}{1 + \frac{10^{2\mathbf{u}+2 \log u_L}}{u_d^2}} \end{aligned} \quad (43)$$

From Eq. (28) and introducing the notation \mathbf{d} of Eq. (21), one has:

$$u_d = 10^{\frac{\mathbf{d} + \log(y_L) - b}{a}} = 10^{\frac{\mathbf{d}}{a} + \log(u_L)} \quad (44)$$

given Eq. (22).

Changing u_d by Eq. (44) in Eq. (43):

$$K = \frac{-1}{1 + 10^{2u+2\log u_L - 2\frac{d}{a} - 2\log(u_L)}} \quad (45)$$

$$= \frac{-1}{1 + 10^{2u - 2\frac{d}{a}}}$$

□

References

- [1] Abul, O., Bonchi, F., Nanni, M., 2008. Never walk alone: Uncertainty for anonymity in moving objects databases, in: 24th International Conference on Data Engineering, IEEE, pp. 376–385.
- [2] Ağır, B., Huguenin, K., Hengartner, U., Hubaux, J.P., 2016. On the privacy implications of location semantics. *Proceedings on Privacy Enhancing Technologies* 2016.
- [3] Agir, B., Papaioannou, T.G., Narendula, R., Aberer, K., Hubaux, J.P., 2014. User-side adaptive protection of location privacy in participatory sensing. *GeoInformatica* 18, 165–191.
- [4] Al-Dhubbani, R., Cazalas, J., 2017. Correlation analysis for geo-indistinguishability based continuous lbs queries, in: 2017 2nd International Conference on Anti-Cyber Crimes, IEEE, pp. 203–208.
- [5] Andrés, M.E., Bordenabe, N.E., Chatzikokolakis, K., Palamidessi, C., 2013. Geo-indistinguishability: Differential Privacy for Location-based Systems, in: CCS, ACM, pp. 901–914.
- [6] Åström, K.J., Hägglund, T., 1995. PID controllers: theory, design, and tuning, 2nd Edition. volume 2. The Instrumentation, Systems, and Automation Society.
- [7] Boutet, A., Ben Mokhtar, S., Bouzouina, L., Bonnel, P., Brette, O., Brunie, L., Cunche, M., D’alu, S., Primault, V., Raveneau, P., Rivano, H., Stanica, R., 2017. PRIVA’MOV: Analysing Human Mobility Through Multi-Sensor Datasets. *NetMob*.
- [8] Brenner, H., Nissim, K., 2014. Impossibility of differentially private universally optimal mechanisms. *SIAM Journal on Computing* 43, 1513–1540.
- [9] Capanema, C.G.S., Silva, F.A., Silva, T.R.B., Loureiro, A.A., 2021. Dcluster: Geospatial analytics with poi identification. *Journal of Information and Data Management* 12.
- [10] Cerf, S., Bouchenak, S., Robu, B., Marchand, N., Primault, V., Mokhtar, S.B., Boutet, A., Chen, L.Y., 2021. Automatic privacy and utility preservation for mobility data: A nonlinear model-based approach. *IEEE Transactions on Dependable and Secure Computing* 18, 269–282. doi:10.1109/TDSC.2018.2884470.
- [11] Cerf, S., Robu, B., Marchand, N., Bouchenak, S., 2023. Privacy protection control for mobile apps users. <https://gitlab.com/cerso/control-of-location-privacy>.
- [12] Chatzikokolakis, K., Palamidessi, C., Stronati, M., 2015. Constructing elastic distinguishability metrics for location privacy, in: PETS, pp. 156–170.
- [13] Chow, C.Y., Mokbel, M.F., Liu, X., 2006. A peer-to-peer spatial cloaking algorithm for anonymous location-based service, in: Proceedings of the 14th annual ACM international symposium on Advances in geographic information systems, pp. 171–178.
- [14] Duckham, M., Kulik, L., 2005. A formal model of obfuscation and negotiation for location privacy, in: International conference on pervasive computing, Springer, pp. 152–170.
- [15] Dwork, C., 2006. Differential Privacy, in: Automata, Languages and Programming. Springer Berlin Heidelberg, volume 4052 of *Lecture Notes in Computer Science*, pp. 1–12.
- [16] Farokhi, F., Eshfahani, P.M., 2018. Security versus privacy, in: IEEE Conference on Decision and Control (CDC), pp. 7101–7106. doi:10.1109/CDC.2018.8619460.
- [17] Farokhi, F., Milosevic, J., Sandberg, H., 2016. Optimal state estimation with measurements corrupted by laplace noise, in: IEEE 55th Conference on Decision and Control (CDC), pp. 302–307. doi:10.1109/CDC.2016.7798286.
- [18] Filieri, A., Hoffmann, H., Maggio, M., 2014. Automated design of self-adaptive software with control-theoretical formal guarantees, in: Proceedings of the 36th International Conference on Software Engineering, Association for Computing Machinery, New York, NY, USA, p. 299–310. doi:10.1145/2568225.2568272.
- [19] Foursquare, . How to get POI data right. URL: <https://foursquare.com/article/how-to-get-poi-data-right/>.
- [20] Franceschi-Bicchierai, L., 2015. Redditor cracks anonymous data trove to pinpoint muslim cab drivers. <http://mashable.com/2015/01/28/redditor-muslim-cab-drivers/>.
- [21] Gambs, S., Killijian, M.O., Cortez, M.N.d.P., 2013. De-anonymization Attack on Geolocated Data. 2013 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications , 789–797doi:10.1109/TrustCom.2013.96.
- [22] Gambs, S., Killijian, M.O., del Prado Cortez, M.N., 2011. Show Me How You Move and I Will Tell You Who You Are. *Transactions on Data Privacy* 4, 103–126.
- [23] Gambs, S., Killijian, M.O., del Prado Cortez, M.N., 2012. Next place prediction using mobility markov chains, in: Proceedings of the First Workshop on Measurement, Privacy, and Mobility, ACM, p. 3.
- [24] Google Play, . Travel & Local - Android Apps on Google Play. URL: https://play.google.com/store/apps/category/TRAVEL_AND_LOCAL?hl=en&gl=US.
- [25] Han, S., Topcu, U., Pappas, G.J., 2016. Differentially private distributed constrained optimization. *IEEE Transactions on Automatic Control* 62, 50–64.
- [26] Hariharan, R., Toyama, K., 2004. Project lachesis: parsing and modeling location histories, in: International Conference on Geographic Information Science, Springer, pp. 106–124.
- [27] Jiang, H., Li, J., Zhao, P., Zeng, F., Xiao, Z., Iyengar, A., 2021. Location privacy-preserving mechanisms in location-based services: A comprehensive survey. *ACM Computing Surveys* 54, 1–36.
- [28] Kang, H., Zhang, S., Jia, Q., 2019. A method for time-series location data publication based on differential privacy. *Wuhan University Journal of Natural Sciences* 24, 107–115.
- [29] Kawano, Y., Cao, M., 2018. Revisit input observability: A new approach to attack detection and privacy preservation, in: 2018 IEEE Conference on Decision and Control (CDC), pp. 7095–7100.
- [30] Kephart, J.O., Chess, D.M., 2003. The vision of autonomic computing. *Computer* , 41–50.
- [31] Koufogiannis, F., Pappas, G.J., 2016. Location-dependent privacy, in: 55th Conference on Decision and Control, IEEE, pp. 7586–7591.
- [32] Krumm, J., 2007. Inference attacks on location tracks, in: International Conference on Pervasive Computing, Springer, pp. 127–143.
- [33] Lau, B.P.L., Hasala, M.S., Kadaba, V.S., Thirunavukarasu, B., Yuen, C., Yuen, B., Nayak, R., 2017. Extracting point of interest and classifying environment for low sampling crowd sensing smartphone sensor data, in: 2017 IEEE International Conference on Pervasive Computing and Communications Workshops, IEEE, pp. 201–206.
- [34] Lemos, R.d., Garlan, D., Ghezzi, C., Giese, H., Andersson, J., Litoiu, M., Schmerl, B., Weyns, D., Baresi, L., Bencomo, N., et al., 2017. Software engineering for self-adaptive systems: Research challenges in the provision of assurances, in: Software Engineering for Self-Adaptive Systems III. Assurances. Springer, pp. 3–30.
- [35] Liao, L., Fox, D., Kautz, H., 2007. Extracting places and activities from gps traces using hierarchical conditional random fields. *The International Journal of Robotics Research* 26, 119–134.
- [36] Maouche, M., Ben Mokhtar, S., Bouchenak, S., 2018. Hmc: Robust privacy protection of mobility data against multiple re-identification attacks. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1–25.
- [37] Oya, S., Troncoso, C., Pérez-González, F., 2017. Back to the drawing board: Revisiting the design of optimal location privacy-preserving mechanisms, in: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1959–1972.

- [38] Palia, A., Tandon, R., 2018. Optimizing noise level for perturbing geo-location data, in: Future of Information and Communication Conference, Springer. pp. 63–73.
- [39] Parekh, S., Gandhi, N., Hellerstein, J., Tilbury, D., Jayram, T., Bigus, J., . Using control theory to achieve service level objectives in performance management., in: IEEE International Symposium on Integrated Network Management, pp. 14–18.
- [40] Piorkowski, M., Sarafijanovic-Djukic, N., Grossglauser, M., 2009. CRAWDAD dataset epfl/mobility (v. 2009-02-24). Downloaded from <http://crawdad.org/epfl/mobility/20090224>. doi:10.15783/C7J010.
- [41] Primault, V., Ben Mokhtar, S., Lauradoux, C., Brunie, L., 2015. Time distortion anonymization for the publication of mobility data with high utility, in: TrustCom, pp. 539–546.
- [42] Primault, V., Boutet, A., Mokhtar, S.B., Brunie, L., 2016. Adaptive location privacy with alp, in: Symposium on Reliable Distributed Systems, IEEE. pp. 269–278.
- [43] Primault, V., Boutet, A., Mokhtar, S.B., Brunie, L., 2018. The long road to computational location privacy: A survey. IEEE Communications Surveys & Tutorials .
- [44] Shevtsov, S., Berekmeri, M., Weyns, D., Maggio, M., 2018. Control-theoretical software adaptation: A systematic literature review. IEEE Transactions on Software Engineering 44, 784–810.
- [45] Shokri, R., Theodorakopoulos, G., Le Boudec, J.Y., Hubaux, J.P., 2011. Quantifying location privacy, in: 2011 IEEE symposium on security and privacy, IEEE. pp. 247–262.
- [46] Srivastava, V., Naik, V., Gupta, A., 2014. Privacy breach of social relation from location based mobile applications, in: International Conference on Contemporary Computing, IEEE. pp. 324–328.
- [47] Sweeney, L., 2002. k-anonymity: A model for protecting privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems 10, 557–570.
- [48] The Guardian, 2018. Fitness tracking app strava gives away location of secret us army bases .
- [49] Vrancic, D., 1996. Design of anti-windup and bumpless transfer protection. Ph.D. thesis. University of Ljubljana, J. Stefan Institute.
- [50] Wang, Y., Huang, Z., Mitra, S., Dullerud, G.E., 2014. Entropy-minimizing mechanism for differential privacy of discrete-time linear feedback systems., in: CDC, pp. 2130–2135.
- [51] Wang, Y., Xu, D., He, X., Zhang, C., Li, F., Xu, B., 2012. L2p2: Location-aware location privacy protection for location-based services, in: INFOCOM, 2012 Proceedings IEEE, IEEE. pp. 1996–2004.
- [52] Wills, A., Schön, T.B., Ljung, L., Ninness, B., 2013. Identification of hammerstein–wiener models. Automatica 49, 70–81.
- [53] Xiao, Y., Xiong, L., 2015. Protecting locations with differential privacy under temporal correlations, in: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, ACM. pp. 1298–1309.
- [54] Zhang, W., Li, M., Tandon, R., Li, H., 2018. Online location trace privacy: An information theoretic approach. IEEE Transactions on Information Forensics and Security 14, 235–250.