Privacy enhancing techniques for the protection of mobility patterns in LBS: research issues and trends

Maria Luisa Damiani

March 31, 2012

1. Introduction

Computing and online services are increasingly being consumed through mobile devices, including smartphones and tablets. Indeed, more than half of the world population now owns mobile phones, which are capable of running applications in ways that involve the collection, use and sharing of location data[1]. Location-based services (LBS) have become an integral part of users' experiences and an increasingly important market. They deliver to users targeted, relevant and highly convenient information, such as up-to-the-minute traffic reports; the location of the nearest petrol stations, hospitals, or banks; as well as targeted advertisements and coupons for services located in a consumer's immediate range. However, the significant advantages associated with LBS come at a price to users' privacy. While sporadic positions of a mobile device may not be particularly sensitive, the historical trail of past locations, i.e. the user's *trajectory*, can reveal much about a user's behavior. In fact, positioning systems allow constant monitoring of the users' position, both indoors and outdoors; moreover techniques for mobility patterns discovery are increasingly deployed in real applications to summarize users' movement and extract behavioral information, e.g. users' activities, from trajectory data.

Location PETs are privacy enhancing techniques conceived to protect position information from privacy violations in on-line applications. Related literature is rich in location PETs offering solutions to diverse privacy requirements for different typologies of on-line services[2], such as policy-based location PETs and techniques for the protection of *identity privacy* and *location privacy[3].* In this paper we argue that conventional location PETs do not have the ability to prevent the extraction of behavioral information from trajectory data collected through LBS, mostly because these techniques ignore the context in which users are located. *Position context* plays a fundamental role in the understanding of the users' behavior in pervasive settings[4]. In particular it can reveal what the person is doing, e.g. a person staying in a clinic for a few days is very likely a person who has been hospitalized, while two persons frequenting the same fitness club in the same period, very likely know each other. Preventing the extraction of behavioral information calls for techniques capable of recognizing mobility patterns based on the geographical, temporal and social context.

To support this argument, in what follows we bring examples of behavioral information which can be extracted from trajectory data. Next we discuss the limitations of conventional classes of location PETs. We also consider the aspect of privacy usability[5], because this is a major requirement for the effective deployment of location PETs, where "usability relates not only to understanding what taking a particular action means in the context of a particular interaction, but also to whether the user understands the

[1] Ericsson. Market and data report. Nov. 2011. http://www.ericsson.com/res/docs/2012/tmd_report_feb_web.pdf
[2] John Krumm. "A survey of computational location privacy" 2009
[3] Christian Jensen et al. "Location Privacy Techniques in Client-Server Architectures". 2009
[4] Pankaj Mehra. "Context-Aware Computing: Beyond Search and Location-Based Services, 2012
[5] Giovanni Iachello et al. "End-User Privacy in Human-Computer Interaction". 2007

implications of his or her choices in a broader context"[6]. Finally we introduce recent research on *semantic location privacy* which aims at protecting the *places* (or semantic location) in which users stay, e.g. hospital. These techniques are a first step in the direction of more effective protection of user's behavior.

The rest of the paper is organized in three sections: section 2 introduces the application context and privacy requirements; section 3 overviews the features of four classes of location PETs, including the aforementioned "conventional" techniques and semantic location privacy techniques; the conclusive section 4 covers additional privacy requirements originating from the recent diffusion of positioning services offered by third party providers and reports some final considerations.

2. Technological and application context

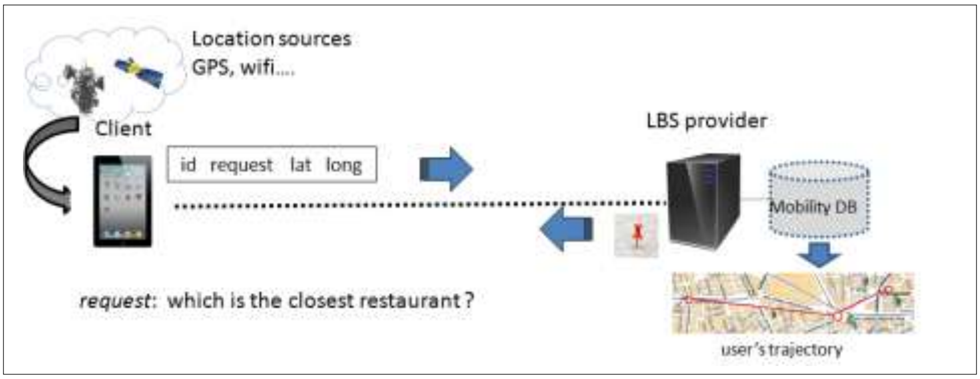

Figure 1 The conventional architecture of a LBS application

Figure 1 illustrates the two main components of a conventional LBS application: (a) a set of location-aware mobile devices, acting as *clients*, i.e. requesters of the information service, which acquire their (accurate) position through a GPS receiver or some other trustworthy location source; b) The LBS *provider* which acts as *server*, i.e. it responds to the requests of service by providing geo-referenced information tailored to the client's position. The requester of the service specifies its identifier, e.g. IP address, the service, e.g. a query, and the position coordinates, e.g. latitude and longitude. The LBS provider stores the position information along with supplementary information in a *mobility database*. A sequence of time-stamped positions forms a user's trajectory.

2.1 Extracting behavioral information from trajectory data in LBS: an example

In certain applications users are allowed to inspect the content of the mobility database. For example, the users of the location sharing service Google Latitude[7] can use the *Location History* functionality to store, view, and manage their past Latitude locations. Figure 2.(b) illustrates the trajectory of a volunteer user running the Latitude application on a smartphone. Following common usage, the device is permanently connected to Internet, therefore the user's position is constantly monitored by the application running in background. The trajectory, reporting the movement during one week in Milan, is displayed as sequence of segments, each connecting two consecutive positions. A dashboard allows inspection of the content, for example, by regulating the time-bar (at the bottom of the picture) one can find where a person was located at a precise instant and how long the person stayed in that position. Moreover, as the trajectory is drawn onto a detailed map, the places that the user visits can be easily identified.

[6] Security Steering Committee on the Usability and Privacy of Computer Systems; National Research Council. "Overview of Security, Privacy, and Usability". 2010

[7] http://www.google.com/latitude

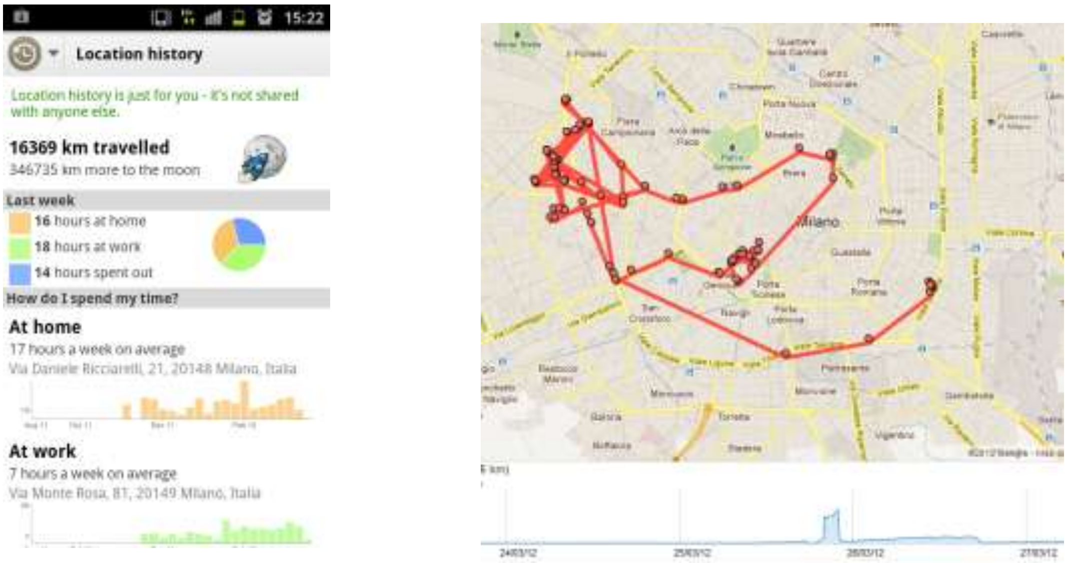

15:22

Location history

Location history is just for you - it's not shared with anyone else.

16369 km travelled
346735 km more to the moon

Last week
16 hours at home
18 hours at work
14 hours spent out

How do I spend my time?

At home
17 hours a week on average
Via Daniele Ricciarelli, 21, 20148 Milano, Italia

At work
7 hours a week on average
Via Monte Rosa, 81, 20149 Milano, Italia

Figure 2. Google Location History: movement statistics (a) and the trace of an individual (b)

More interesting is Figure 2.(a) which illustrates the statistics that the system provides on the user's activities, in particular the time spent at home, at work and outside. Note that the patterns "home", "work" are inferred from the system based on the movement information. For example, the inactivity periods during night hours can reveal where the user lives, while frequent movements from home to some other place at certain hours can disclose where the user works.

While it can be seen that the accuracy of the extracted information is low (e.g. the time spent at home is unrealistic) this example clearly shows the potential of the techniques for mobility patterns discovery. It is also foreseeable that information accuracy will rapidly increase in the near future under the push of on-going research on mobility pattern discovery and representation. On-going research includes, for example, *trajectory pattern mining*[8] which aims at identifying the regions that people usually frequent, how much time is spent in each of those regions and the preferred order in which those regions are visited; *mining of points of interest*[9] i.e. extraction of places that are significantly frequented; *semantic trajectories*[10] which allow the representation of behavioral information in a machine-readable form.

2.2 Mobility patterns

Mobility patterns reveal what people do, i.e. behavioral information. For example, people spend different amount of time in a location depending on what they do there, e.g. a user staying in a night-club at nightly hours is likely a customer of the nightspot. This pattern is called *staying duration* in Lee et al.[11] Other interesting examples of patterns, besides the home-work pattern seen in the previous example, are reported in Opinion 13/20111 by the Article 29 Working Party[12]. In particular, patterns may include data derived from the movement patterns of friends as well as "special categories of data", such as visits to hospitals and religious places, presence at political demonstrations or presence at other specific locations revealing data about, for example, sex life. In all these examples, the extraction of behavioral information is leveraged by the intertwining of trajectory data with contextual information such as geographical places, time, frequency, duration of staying, and the social context

[8] Fosca Giannotti et al. "Trajectory pattern mining". 2007
[9] Xin Cao et al. "Mining significant semantic locations from GPS data". 2010
[10] Stefano Spaccapietra et al. "A conceptual view on trajectories".2008
[11] Byoungyoung Lee et al. "Protecting Location Privacy Using Location Semantics". 2011
[12] Article29 Data Protection Working Party. Opinion 13/2011. 2011

3 Privacy enhancing techniques for the protection of position in LBS

The bulk of research on privacy of position data took off with the emergence of mobile applications based on stored people's tracks[13],[14] early past decade. Current location PETs can be grouped in two broad classes of solutions. The first class of techniques are commonly referred to as *policy-based*. A *policy* consists of a set of user-defined privacy preferences or *rules* typically enforced by the trustworthy LBS provider upon the request of service. We refer to the second class of solutions as *inference-prevention* techniques. These techniques basically aim at preventing the LBS provider from drawing sensitive information from exact positional data. Note that in this case the LBS provider is considered not fully trustworthy, for example it can cooperative and curious. Taking inspiration from the classification proposed by Jensen et al.[15], we further categorize inference-prevention techniques in the following classes:

- *Identity privacy* techniques attempt to forestall the re-identification of users (deprived of their real identity) in LBSs providing anonymous services
- *Location privacy* techniques apply to forestall the transmission of *exact* users' positions to the LBS provider. Knowing precisely the positions in which individuals are located (or not located) jeopardizes their privacy and physical safety.
- *Semantic location* privacy techniques aim at preventing the disclosure of the places in which users stay because those locations can reveal sensitive data and behavioral information.

The whole taxonomy is shown in Figure 3.

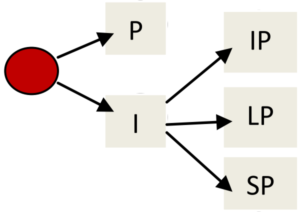

P – Policy-based techniques
I – Inference prevention techniques
IP – Identity privacy
LP – Location privacy
SP – Semantic location privacy

Figure 3 A taxonomy of location PETs

In the next, we examine these four classes of techniques, i.e. policy-based techniques and the three inference prevention techniques. In order to keep the paper focused, we choose not to use any formal privacy and utility metric[16], while this analysis is postponed for future work.

3.1 Policy-based techniques

Policy-based techniques are probably the most popular solutions for privacy in LBS, conceptually simple, in line with common practices in law, and endorsed by standardization bodies such as IETF Geopriv[17]. These techniques allow users to specify which position is to be disclosed to whom and when, through a set of machine-readable and enforceable privacy rules. *Machine-readable* means that rules are encoded using a computer language (i.e. a *policy-specification language*) instead of being expressed in natural language; *enforceable* means that those rules can be checked by an automated system, on behalf of the user. These techniques have their roots in security, in particular in access control policies, and in the bulk of work developed at the end of nineties for privacy protection in e-commerce applications, i.e. P3P[18].

As an example, consider the case in which the user John wants to share his position with acquaintances through a location sharing service constantly monitoring the user's position. Because acquaintances include colleagues, relatives, and friends, John chooses to specify different rules, one for each category. A rule can state for example that John's position can be revealed to colleagues Bob and Mary exclusively when John is

[13] Mark Gruteser et al. Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. 2003
[14] Alastair Beresford et al. "Location Privacy in Pervasive Computing". 2003
[15] Christian Jensen et al. "Location Privacy Techniques in Client-Server Architectures". 2009
[16] Reza Shokri et al. "Quantifying Location Privacy". 2011
[17] IETF. "An Architecture for Location and Location Privacy in Internet Application". 2011
[18] Lorrie Cranor. "*P3P*: Making Privacy Policies More Useful". 2003

at work and during working-time. The set of rules forms the John's privacy policy. For example, this technique is used in the Locaccino location sharing service[19]. In particular, the subscribers of this service can specify privacy rules encompassing both temporal conditions and spatial conditions, i.e. the periods and the regions within which the position can be disclosed or hidden to acquaintances. These rules are finally enforced by the Locaccino server.

Discussion: policy-based techniques do not prevent the extraction of mobility patterns because the LBS provider is generally aware of the positions of all clients and thus can record users' trajectories at the finer level of detail. Therefore, if the LBS provider is untrustworthy, the user's privacy is at stake. However policy specification languages have a peculiar feature, i.e. the capability of expressing conditions on contextual variables. The degree of usability of these languages is generally assessed by involving users in the experimentation. For example, Tsai et al. report the positive feedback of a group of selected users requested to use solely time-based privacy rules such as: "Show my location between 9 am and 6 pm on Mondays and Wednesdays"[20].

3.2 Identity privacy techniques

Identity privacy techniques are conceived to forestall the re-identification of seemingly anonymous users, based on position information. For example, consider the case in which an LBS is offered to the members of a community potentially subject to discrimination, e.g. the gay community, and assume users to interact with the system through pseudo-identifiers. Unfortunately simply stripping off users' identifiers is not sufficient to ensure anonymity, because the LBS provider can draw users' identities from trajectory information, e.g. if a user requests the service from a certain place early in the morning, it is likely that such a place is his or her home and thus the user can be easily re-identified using a white pages service [21]. While we refer the reader to Chow et al. for a recent survey on trajectory privacy [22], we limit ourselves to consider an exemplifying paradigm, i.e. *location k-anonymity*.

Given a population of users, location k-anonymity postulates the following requirement, that the user's position disclosed to the LBS provider must be indistinguishable from the position of at least k-1 other users. In practice, the exact user's position must be replaced by a coarser position, i.e. a *cloaked* region, large enough to contain the position of *k-1* other users located nearby at the time the on-line service is requested. Accordingly, the LBS provider cannot identify the requester of the service based exclusively on the position information. This situation is exemplified in Figure 4. For k=10, the position of the single individual is replaced by a larger region (i.e. a cloaked region) containing 10 persons. If the on-line service is requested from this region, the maximum probability of identifying the requester is 1/10. Another prominent feature of this privacy mechanism is that it typically requires a dedicated trusted middleware, the *location anonymizer*, in between the clients and the LBS provider. The location anonymizer is aware of the position of all the clients, intercepts the individual's requests, replaces the user's identifier with a pseudo-identifier and finally replaces the true position with the dynamically generated cloaked region. One representative solution of this class is the Casper system[23] (Figure 5). Casper consists of the location anonymizer and the *privacy-aware query processor*, a software component which runs on the server (i.e. the LBS provider), and which resolves user's requests with respect to a position which is not a point as usual, but a region and which returns a set of candidate answers. Although alternative architectures have been proposed[24], the practical deployment of location k-anonymity in real applications looks complex and costly.

[19] Eran Toch et al. "Locaccino: a privacy centric location sharing appplication". 2010

[20] Tsai, Janice et al. Who's viewed you?: the impact of feedback in a mobile location-sharing application. 2009.

[21] Mark Gruteser et al. "Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking". 2003

[22] Chi-Yin. Chow et al."Trajectory privacy in location-based services and data publication". 2011

[23] Mohamed Mokbel et al. The new Casper: query processing for location services without compromising privacy. 2006

[24] Gabriel Ghinita et al."MobiHide: A Mobile Peer-to-Peer System for Anonymous Location-Based Queries". 2007

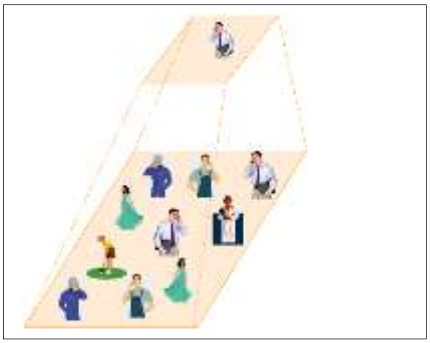

Figure 4 A cloaked region for K=10

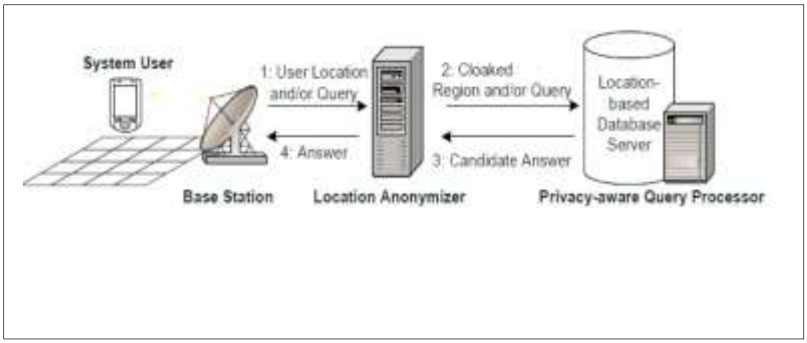
System User
1: User Location and/or Query
4: Answer
2: Cloaked Region and/or Query
3: Candidate Answer
Location-based Database Server
Base Station
Location Anonymizer
Privacy-aware Query Processor

Figure 5 The Casper system

Discussion: location k-anonymity techniques do not forestall the extraction of mobility patterns from trajectory data (even though trajectories have a coarse granularity), because the position context is ignored. For example, cloaked regions are generated independently from the geographical setting. Consequently, if a cloaked region falls inside the area covered by a hospital, one can infer that the k users grouped in the region suffer from health concerns. Hence, if users are re-identified, there is a privacy leak. In essence, location k-anonymity only serves to protect the association between users and service requests. Another consideration regards usability. It is difficult to gauge which size of k is minimally necessary or sufficient [25]. The higher the value of k, the higher the level of protection but also the loss of position accuracy (and thus of quality of service, *QoS*), where the position accuracy varies in time and space based on the distribution of people.

3.3 Location privacy techniques

Location privacy techniques aim at preventing the disclosure of exact users' position in the context of LBSs providing non-anonymous services, for example geo-social networks[26]. These techniques communicate to the LBS provider a location other than the exact position. In particular, the disclosed position can be fake, cloaked or can be transmitted using some cryptographic protocol.

- A *fake position* is a position deliberately represented with a wrong value. Privacy is achieved from the fact that the reported position is false. The accuracy and the amount of privacy mainly depend on how far the reported location is from the exact location. For example, the client requesting a service, e.g. "where is the closest restaurant" can transmit to the LBS provider a fake position and then properly fil
- ter out candidate answers.[27]
- A*n obfuscated position* is another term for cloaked region including the exact user's location. Therefore the LBS provider does know that the user is located in the cloaked region, but has no clue where exactly the user is located. A popular obfuscation method[28], also used in commercial platforms[29], replaces the actual position with a predefined region chosen in a taxonomy of locations at different granularities e.g. street, zip code area, city. Unfortunately predefined locations can be too broad to ensure an appropriate quality of service, say a zip code region covering an area of few squared kilometers, or conversely too small to provide privacy guarantees, say a short street. Another simple method obfuscates the position with a circle of user-defined radius and random center containing the actual position[30]. In more complex solutions, the size of

[25] Mark Gruteser et al., 2003, see note 19

[26] Carmen Ruiz Vicente et al. "Location-Related Privacy in Geo-Social Networks". 2011

[27] Man Lung Yiu et al. "SpaceTwist: Managing the Trade-Offs Among Location Privacy, Query Performance, and Query Accuracy in Mobile Services". 2008

[28] IETF, 2011, see note 16

[29] http://fireeagle.yahoo.net/

[30] Claudio Ardagna et al. "Location privacy protection through obfuscation-based techniques". 2006

the cloaked region is the result of the trade-off between privacy and QoS[31] while the transmission of the position can be also delayed a while to cloak the temporal dimension[32].

- Cryptographic protocols define techniques for the secure collaboration of different parties. An example of *cryptographic protocol* used in LBS is PIR (Private Information Retrieval). This technique allows users to issue a query without disclosing to the LBS provider the information which is requested as well as the information being returned[33]. In this sense this technique protects both the identity and the location. The method ensures the maximum privacy. However it incurs high computational costs and can be only applied to certain categories of queries, e.g. the retrieval of stationary objects (i.e. non-mobile objects).

One specific problem that may rise when the position is obfuscated by a coarse region is that consecutive positions in the user's trajectory are correlated, i.e. the presence in one region constrains the position in the subsequent regions. This information can be exploited to prune the cloaked regions and more precisely delimitate the user's position. To prevent this inference when the maximum speed of the user is known (e.g., the user can be a pedestrian, a car driver, a cyclist and so on) and the movement is frequently sampled, an approach is to modify the position in space and time before it is released [34].

Discussion: in general, location privacy techniques are not able to prevent the extraction of mobility patterns. The solutions based on obfuscation and fake positions have the same limitations discussed in the previous section, i.e. lack of context awareness, while the deployment of cryptographic protocols in LBS is somewhat limited to specific situations or applications. As concerns the aspect of usability, obfuscation techniques are the simplest but not necessarily usable solutions. For example, what is the loss of QoS if the position is disclosed at the level of zip code area, instead of street? The lack of suitable metrics makes it difficult understanding the implications of certain choices.

3.4 Semantic location privacy techniques: a first step towards the protection of behavioral information

Semantic location privacy techniques attempt to prevent LBS providers from identifying the semantic locations in which users stay.[35,36] For example, one of the motivating observations is that the sensitivity of positions may vary depending on the nature of places, e.g. the position of a user staying in an oncological clinic is likely "more sensitive" than the position of an user walking along a street. Indiscriminately treating every position by imposing the maximum level of privacy for each position would compromises QoS. A more flexible solution is to protect only those positions which are perceived as sensitive, while the others that are not sensitive are disclosed with no change. In this way the loss of QoS can be limited. This form of obfuscation is called *semantic location cloaking.*

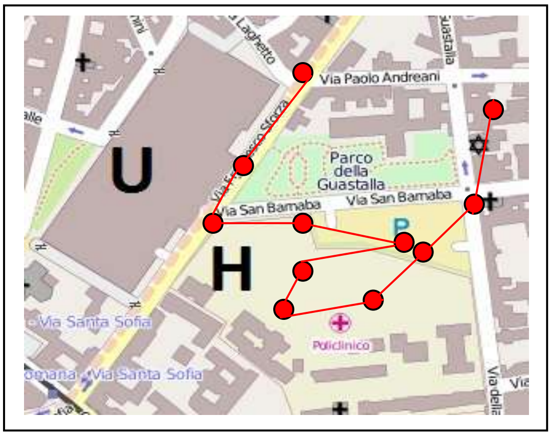

Figure 6. Urban setting and Bob' s route. The map is drawn from http://www.openstreetmap.org

[31] Marc Duckham et al. "A Formal Model of Obfuscation and Negotiation for Location Privacy". 2005

[32] Reynold Cheng et al. "Preserving user location privacy in mobile data management infrastructures". 2006

[33] Gabriel Ghinita et al. " Private queries in location based services: anonymizers are not necessary". 2008

[34] Gabriel Ghinita et al. "Preventing velocity-based linkage attacks in location-aware applications". 2009

[35] Byoungyoung Lee et al. "Protecting Location Privacy Using Location Semantics". 2011

[36] Maria Luisa Damiani et al. "Fine-grained cloaking of sensitive positions in location sharing applications". 2011

As an example, consider the urban setting in Figure 6. The map shows a number of places in Milan: the premises of the Policlinico hospital, the University of Milan, a few religious buildings, various private buildings, and the road network. Assume that the user Bob connects to a location sharing service through a smartphone. Bob is driving his car when in the proximity of the Policlinico hospital, Bob stops in a parking area and steps onto the hospital premises where he remains for a few hours for a medical visit, before again taking the car to reach his friends in a pub in downtown. During this time, Bob's position is continuously reported to the LBS provider as well as his friends, therefore the places in which Bob stops are known, including those that Bob consider sensitive, e.g. the hospital. Simply disconnecting from the service would prevent Bob from being in touch with his friends, unless suspending and then resuming the service which would create considerable burden to Bob. The issue is how not to reveal to the LBS provider that the user certainly stays in a certain place, without giving up the service.

To illustrate the technical issue posed by this problem, consider first a naive solution. Assume a user in position *p*. Upon a request of service, the main steps of the privacy enforcement process are:

1. The client checks whether *p* is within one of the places considered sensitive (assume there is a precompiled list of sensitive places, e.g. hospitals, religious buildings and a map on the client)
2. If this is the case, generate a cloaked region containing the actual position
3. Otherwise, if the user is not in a sensitive place, release the actual position

It is easy to see that, if the LBS provider is aware of the protection strategy, it can promptly infer from the fact that Alice is in a cloaked region that she is certainly inside a sensitive location. Moreover, if the party has clues about the sensitive locations, she can more precisely localize Alice inside the cloaked region. As a result the protection mechanism fails. In previous work[37] we argued that a sound cloaking strategy should guarantee:

- **Semantic diversity**. The user's position cannot be blurred exclusively when the user is inside a sensitive place, but also when he or she is outside. That way, the place in which the user is located remains uncertain. A cloaked region thus must include places of diverse types.
- **Independence** of the position cloaking method from the user's position. This condition prevents the discovery of the correlation between the cloaked region and the true position, which could be exploited to infer where the user is located.

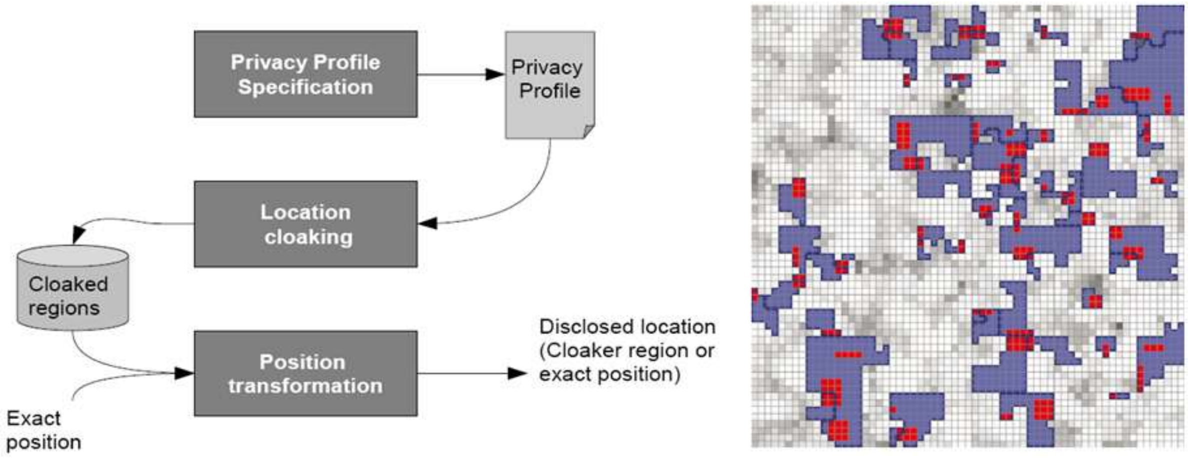

Figure 7 Probe system: (a) the workflow; (b) obfuscated map: the blue polygons represent cloaked regions, the red rectangles sensitive places, the grey background the distribution of population in space

These guidelines have been embodied in the privacy-preserving framework called Probe (Privacy-aware Obfuscation Environment)[38]. Figure 7.a illustrates the workflow of the privacy enforcement process. Users first specify in a privacy profile which categories of points of interest are sensitive (selecting for example from a pre-defined list, e.g. hospitals, religious buildings and so on) along with the degree of privacy desired

[37] See note 36.

[38] Maria Luisa Damiani et al. The PROBE frame work for the protection of sensitive positions". 2010

for each of those categories. For example a privacy degree of 0.1 assigned to hospitals means that the (posterior) probability of locating the user inside a hospital must be less than 0.1. Next, coarse regions are generated satisfying the privacy preferences, independently from the user's position, in order to prevent possible inferences on their reciprocal positions. A sample set of cloaked regions is shown in Figure 7.b. Finally, at runtime if the user's position falls inside one of the coarse regions, that region is delivered instead of the exact position. Recent results extend these techniques to the case in which users' movement is confined to road network[39]. In this case the cloaked region takes the form of a subgraph of a semantically annotated graph representing the urban setting.

Discussion. The concern for semantic location privacy is recent and thus many research issues are still open. For example, an issue is how to intertwine the geographical context with the temporal and social dimension; another problem regards the protection of interrelated places, e.g. the home-work pattern. As concerns usability, no study has been carried out on this aspect. However, in the specific case of the PROBE system, users can specify their privacy preferences in a privacy profile using an intuitive and conceptually founded privacy metric. Moreover an additional metric is defined, the utility metric, providing a measure of the spatial accuracy of the cloaked regions. Unlike more traditional obfuscation techniques, the utility measure can be computed prior to any service request. In this way users can tune and balance the amount of privacy with QoS.

4 Open issues and conclusions

4.1 Towards a more complex LBS model

All the location PETs that we have considered so far, including the most recent techniques, rely on the assumption that the position information is obtained from some trusted source, such as GPS. Indeed, LBSs are rapidly evolving towards novel architectures in which the position information can be pervasively offered by third-party location providers (LP). The location service is offered on a free basis provided that users reveal contextual information, e.g. Wi-Fi networks nearby. Currently, LPs include all major IT players, such as Google and Apple. We postulate that, in the same way of LBS providers, LPs are not necessarily trustworthy.

Architecture and problem formulation. Figure 10 depicts the extended LBS model comprehensive of the third party location provider. A usage scenario is as follows. Assume that a user, equipped with a Wi-Fi enabled device and located in a metropolitan area (with high density of Wi-Fi networks), requests a LBS. The application running on the client handles this request as follows:

- It first determines the position of the device. Since the user is inside a building and thus the GPS signal is not available (or the GPS receiver is not installed), the position is requested from the third party LP. To obtain the position, the client transmits to the LP the set of Wi-Fi access points (APs) and/or the cell towers in proximity of the device. Since the user is located in a metropolitan area, the position can be computed with an accuracy of a few tens of meters.
- Once the coordinates are obtained, the application conveys the position along with the requested service to the LBS provider which returns the requested information as usual.

In this scenario, it is obvious that the LP is necessarily aware of the user's position. Moreover, if the client interacts with a unique LP, such LP is aware of any position flowing to the LBS providers. Now consider the case in which the LP is untrustworthy. It should be clear that existing location PET cannot protect the position from the LP which computes it. Therefore the problem is to what extent privacy can be protected without giving up the LBS and compromising the business model (entailing free access to the LP).

[39] Emre Yigitoglu et al. "Privacy-preserving sharing of sensitive semantic locations under road constraints". 2012

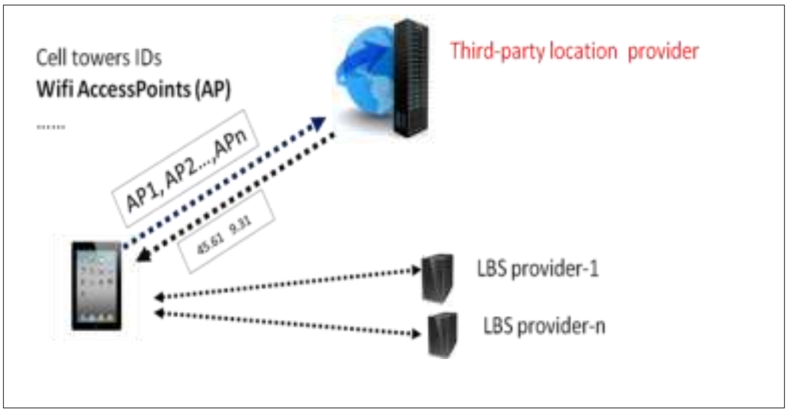

Figure 8 LBS model extended with the third party location provider

Problem analysis. If the client could determine by itself the position with sufficient accuracy both indoors and outdoors, there would not be privacy concerns. Unfortunately providing clients with pervasive geo-location capabilities is costly. We argue that a different approach is to minimize the interaction with the LP. The motivating observation is that the amount of information that the user transmits to the location provider exceeds what is really necessary to determine the users' position. For example every time a service is requested from a place, e.g. home, the client transmits the same or similar contextual information, e.g. Wi-Fi networks in proximity, even though the position has been already obtained the first time a request has been made from that place. Based on this observation, we envision a solution in which enhanced clients can acquire the capability of recognizing places that have been already visited. This way the position is only requested to the LP when it is strictly necessary. We qualify this geo-location service as *privacy-aware*.

To implement this strategy a possible approach is to confine the protection to a subset of positions, in particular those which can be associated with *private places*[40]. Private place is an abstraction which conceptualizes the intuition that there are some regions of space that belong to the personal sphere, e.g. home. Whenever the user is in a private space, the position is not disclosed to the LP.

Note, however, that this solution does not forestall the disclosure of the position to the LBS provider. Therefore for a comprehensive approach, privacy-aware geo-location and (context-aware) location PETs should be integrated.

4.2 Concluding remarks

We conclude with two summarizing considerations:

1. We have seen that location PETs include a variety of techniques conceived to satisfy different privacy requirements. In general, conventional techniques are not able to prevent the extraction of mobility patterns from trajectory data. We have also outlined the features of a recent stream of research for the protection of presence in places, which attempts to introduce the contextual dimension in privacy. This experience can be extended along several directions, for example to account for the temporal and social dimension of privacy. Another interesting research direction regards the combined use of policy specification languages and inference prevention techniques.
2. The architecture and inner workings of current LBS ecosystem remain opaque and largely unknown to users. For example, users often do not know that while they interact with and authorize a specific online or mobile application (Apps) to determine their location, such an App refers to a LP to obtain the localization service. Like many other privacy and data protection problems, transfers of data to LPs need to be addressed through a combination of legal and technological mechanisms. Technological solutions can provide users even more robust privacy protections than legal rules. However, protecting mobility patterns from location providers and LBS providers especially if both parties are untrustworthy, is a challenge.

[40] Maria Luisa Damiani. "Third party geo-location services: privacy requirements and research issues". 2011

Bibliography

1. Claudio. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, and Pierangela Samarati. “Location privacy protection through obfuscation-based techniques.” In Proceedings of the 21st annual IFIP WG 11.3 working conference on data and applications security. 2007
2. ARTICLE 29 Data Protection Working Party. “Opinion 13/2011 on Geolocation services on smart mobile devices”. http://ec.europa.eu/justice/policies/privacy/index_en.htm, 2011
3. Alastair Beresford, Frank Stajano. “Location Privacy in Pervasive Computing”. IEEE Pervasive Computing, 2, 46-55, 2003
4. Xin Cao, Gao Cong, and Christian S. Jensen. “Mining significant semantic locations from GPS data”. *Proc. VLDB Endow.* 3, 1-2 ,1009-1020, 2010
5. Chi-Yin. Chow and Mohamed. Mokbel, “Trajectory privacy in location-based services and data publication”. SIGKDD Explorations, vol. 13, no. 1, pp. 19–29, 2011.
6. Lorrie Cranor. “P3P: Making Privacy Policies More Useful”. IEEE Security and Privacy, 1(6):50–55, 2003
7. Maria Luisa Damiani. “Third party geo-location services: privacy requirements and research issues”. Transaction on Data Privacy, 4(2): 55-72, 2011
8. Maria Luisa Damiani, Elisa Bertino, Claudio Silvestri. “The PROBE Framework for the Personalized Cloaking of Private Locations”. Transactions on Data Privacy, 3(2):123–148, 2010
9. Maria Luisa. Damiani, Claudio Silvestri, Elisa Bertino. “Fine-grained cloaking of sensitive positions in location sharing applications”. IEEE Pervasive Computing, 10(4): 64-72, 2011
10. Nick Doty, Deirdre Mulligan, Erik Wilde. “Privacy issues of the W3C Geolocation API”. Technical report, UC Berkeley, School of Information, 2010.
11. Matt Duckham, Lars Kulik. “A Formal Model of Obfuscation and Negotiation for Location Privacy”. Pervasive Computing, Springer, 152-170,2005
12. Gabriel Ghinita, Panos Kalnis, Spiros Skiadopoulos: MobiHide: A Mobile Peer-to-Peer System for Anonymous Location-Based Queries. SSTD, 221-238. 2007
13. Gabriel Ghinita, Maria Luisa Damiani, Claudio Silvestri, Elisa Bertino. “Preventing velocity-based linkage attacks in location-aware applications”. In Proc. of the ACM International Conference on Advances in Geographic Information Systems, 2009
14. Gabriel Ghinita, Panos Kalnis, Ali Khoshgozaran, Cyrus Shahabi, and Kian-Lee Tan.” Private queries in location based services: anonymizers are not necessary”. In Proceedings of the ACM SIGMOD international conference on Management of data*.* 2008
15. Fosca Giannotti, Mirco Nanni, Fabio Pinelli, Dino Pedreschi. “Trajectory pattern mining”. ACM SIGKDD International conference on Knowledge discovery and data mining. 2007
16. Mark Gruteser and Dirk Grunwald. “Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking”. In Proc. of the 1st International Conference on Mobile systems, Application and Services.2003
17. Giovanni Iachello, Jason Hong. “End-User Privacy in Human-Computer Interaction,” Foundations and Trends in Human-Computer Interaction, vol. 1, no. 1, pp. 1–137, 2007
18. Christian S. Jensen, Hua Lu, and Man Lung Yiu. “Location Privacy Techniques in Client-Server Architectures”. In Privacy in Location-Based Applications, Claudio Bettini, Sushil Jajodia, Pierangela Samarati, and X. Sean Wang (Eds.). Lecture Notes In Computer Science, Vol. 5599. Springer-Verlag, Berlin, Heidelberg 31-58, 2009
19. John Krumm.” A survey of computational location privacy. Personal and Ubiquitous Computing” 13(6). 391-399 2009
20. Pankaj Mehra. "Context-Aware Computing: Beyond Search and Location-Based Services," IEEE Internet Computing, pp. 12-16, March-April, 2012
21. Mohamed Mokbel Chin-Yin Chow, Walid Aref. “The new Casper: query processing for location services without compromising privacy”. In Proc. of Very Large Database Conference, 2006
22. Security Steering Committee on the Usability and Privacy of Computer Systems; National Research Council. “Overview of Security, Privacy, and Usability”. In: Toward Better Usability, Security, and Privacy of Information Technology: Report of a Workshop. The National Academies Press. 2010
23. Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, Jean-Pierre Hubaux. “Quantifying Location Privacy”. IEEE Symposium on Security and Privacy, 247-262, 2011
24. Carmen Ruiz Vicente, Dario Freni, Claudio Bettini, Christian Jensen. “Location-Related Privacy in Geo-Social Networks”. IEEE Internet Computing, 15(3), pp. 20-27, 2011
25. Janice Tsaj et al. “Who’s viewed you?: the impact of feedback in a mobile location-sharing application”. In Proc. of the 27th International Conference on Human Factors in Computing Systems(CHI ’09), 2009
26. IETF, Request for Comments 6280. “An Architecture for Location and Location Privacy in Internet Application”, http://tools.ietf.org/html/rfc6280. 2011

27. Eran Toch, Justin Cranshaw, Paul Hankes-Drielsma, Jay Springfield, Patrick Gage Kelley, Lorrie Cranor, Jason Hong, and Norman Sadeh. “Locaccino: a privacy-centric location sharing application”. In Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing. 2010
28. Emre Ygitoglu, Maria Luisa Damiani, Osman Abul, Claudio Silvestri. “Privacy-preserving sharing of sensitive semantic locations under road constraints”. IEEE International Conference on Mobile Data Management, July 2012.
29. Man Lung Yiu, Christian S. Jensen, Xuegang Huang, and Hua Lu. “SpaceTwist: Managing the Trade-Offs Among Location Privacy, Query Performance, and Query Accuracy in Mobile Services”. In Proceedings of the IEEE 24th International Conference on Data Engineering(ICDE '08). 2008