Chapter 4
Privacy Protection in Location-Based Services: A Survey

Claudio Bettini

Abstract Location awareness has enabled efficient and accurate geo-localised Internet services. Mobile apps exploiting these services have changed our way of navigating and searching for resources in geographical space. This chapter provides a classification of location based services (LBS) and illustrates the privacy aspects involved in releasing our location information as part of a service request. It includes a discussion about legal obligations of the LBS provider and about ways to specify personal location privacy preferences. The chapter also provides a systematic survey of the main approaches that have been proposed for protecting the user's privacy while using these services.

4.1 Introduction

A large majority of smartphone users take advantage of apps that provide Location Based Services (LBS), from weather forecast to map based navigation and search of nearby resources.[1]

LBS can be generally described as Internet based services that offer functionalities enabled by the geo-localization of the device issuing the service request. They belong to the family of context-aware services with the timestamped location and/or the user trajectory acting as the context information. The service provisioning often involves performing some spatio-temporal data processing that may consider the relative distance between users, their velocities, the nearby resources, as well as spatial constraints such as road networks, or real time events like traffic conditions. LBS provisioning may also involve data analysis as, for example, discovering the users' frequent trajectories.

[1] Pew Research Center http://www.pewinternet.org/2013/09/12/location-based-services/.

C. Bettini (✉)
Dipartimento di Informatica, University of Milan, Milan, Italy
e-mail: claudio.bettini@unimi.it

A. Gkoulalas-Divanis, C. Bettini (eds.), *Handbook of Mobile Data Privacy*,
https://doi.org/10.1007/978-3-319-98161-1_4

LBS are not exclusive to mobile devices. Indeed a LBS, for example a local weather forecast, can also be invoked from a desktop PC obtaining the user position by IP address geo-coding. However, there is no doubt that their huge popularity is due to mobile users. Indeed the integration of global navigation satellite receivers (e.g., GPS) into even the cheapest smartphone, and the availability of very effective outdoor positioning services that combine GPS with methods based on cellular and WiFi signal, have made LBS very effective and available to every mobile user. Indoor positioning is less consolidated than outdoor positioning, but research advances, exploring the use of technologies like WiFi fingerprinting, BT beacons, UWB, ultrasound and more, are promising to reach a very high precision and new LBS are being deployed taking advantage of indoor as well as integrated outdoor and indoor positioning.

In order to understand the possible privacy threats in using LBS it is important to briefly review the different types of LBS and their main properties.

4.1.1 A Classification of LBS

We can divide LBS into two broad categories according to their sharing model: (a) personal services, and (b) social network services. The former provide information to the user that asked for the service, and typically share the user location information only with the service provider. The latter are intended to share the user location information also with a group of users and to receive information based on other users position and/or the relative distance between users.

4.1.1.1 Personal LBS

Personal LBS can be grouped as follows:

- Navigational services. They typically provide instructions to reach a destination based on the user position.
- Resource discovery services. They provide nearby points of interest (ATM, gas station, shop, ...) in response to a user location.
- Local traffic/news/weather services. They provide local information based on the user position.
- Emergency services. They can send operators to the location from which the request was issued (e.g., road assistance).
- Proximity marketing services. They send ads to a user based on her proximity to shops or even items on sale.
- Location based personalization/advertising services. Different information/ads are sent to users depending on their location.

Note that the last two groups follow a *push* model while the others are usually based on a *pull* model.

4.1.1.2 Social Network LBS

Social network LBS can be grouped as follows:

- Geo-SN posting/check-in services. They allow users to associate a timestamped position to a resource that they share (e.g., a picture, video, text) or simply share their current position (check-in). They may also offer resource discovery services based on the density (or other property) of check-ins.
- Friend finder services. They provide information about the proximity of contacts or other users of a Geo-SN.
- Workforce management services. They allow coordination and optimization of mobile workers and items based on their position (e.g., taxi or packages in logistics).
- Location based games. They engage participants in a game that involves users as well as resources geographical positions.

The above categorization is definitely not exhaustive, but it provides a good coverage of currently available LBS.

4.1.2 Privacy Threats in LBS

In this chapter we formally define a privacy threat as follows.

A privacy threat occurs whenever an unauthorised entity can associate with high probability the identity of an individual with private information about that individual.

In the context of *personal LBS* such a threat occurs when the information contained in one or more requests issued by a given user can be used, possibly associated with external information, to associate the user identity with the private information.

In the context of *social network LBS* the above association can be obtained also from requests or postings performed by individuals different from the one involved in the threat [56].

According to some country regulation (e.g., the EU GDPR) any information that is specifically associated with an individual should be considered private, while other regulations refer to specific types of information. However, most guidelines for privacy risk assessment highlight the risk involved in revealing information on political, sexual, and religious orientation, health, financial assets, or closeness to specific individuals or organizations. LBS services play a role in this context because both identity and private information can be directly or indirectly released through a single or a sequence of LBS requests. LBS requests can reveal, for example, (a) information on the specific location of individuals at specific times, (b) movement patterns (specific routes at specific times and their frequency), (c) requests for sensitive services (closest temple for a specific religious worship),

(d) personal points of interest (home, workplace, frequent visits to specific shops, clubs, or institutions). Moreover the above information can also be used to infer when the individual is not where it is supposed to be (*absence privacy* [56]), where it is likely to be at a given hour of a given day, or when and how frequently the individual met other individuals (*co-location privacy* [56]).

Unauthorised use of this information exposes the user to several types of *privacy violation risks* including unsolicited advertising, discrimination, loss of reputation, family and work related issues (with divorce and getting fired among the outcomes), stalking and even exposure to robberies based on absence privacy violations.

4.1.2.1 Adversaries

In the privacy protection literature, as well as in the following of this chapter, the unauthorised entity that can acquire some data exchanged as part of the LBS and that may pose a privacy threat is often called the *adversary*. The LBS service provider can be considered an adversary if he uses the acquired information in any way different from what the user has agreed upon. An external entity that tries to break into the communication channel or into the service provider IT infrastructure and get LBS requests in transit or stored in a database is another adversary. In social network LBS, a generic user can be considered an adversary, for example when he can access a geo-localised post involving a user without his explicit consent. As it will be clear in Sect. 4.3 many privacy protection techniques require accurate modeling of the adversary in order to provide guarantees about their effectiveness.

4.1.2.2 Online Versus Offline Data Release

The literature on privacy protection for LBS has considered separately the problem of protecting privacy at the time a LBS request is sent to the service, and the problem of privately releasing data from a database where the history of requests from different users has been stored.

The first case is named *online data release* and the service provider, as well as any entity that may get access to the content of the issued requests are considered adversaries. The second case is named *offline data release* and the adversaries are the third parties or any entity that obtain data extracted from the database of requests stored at the service provider. This can be statistical data (e.g., aggregated counts on visited locations) or a set of individual data records. Such data is used, for example, for profiling, mobile advertising, or to refine online marketing strategies.

Online data release is intuitively more challenging, since a mechanism should be applied at the client side or through a trusted proxy server at the time of each request issued to the service. This case is also characterised by a limited knowledge available to this mechanism about the whole set of requests and about the location of other service users; information that may be useful for some protection techniques.

4.1.2.3 Single Versus Multiple Data Release

In both online and offline data release we should distinguish the case of a single data release from the case of multiple data releases. In the case of tabular data (e.g., databases) it is well known that the privacy guarantee that a protection technique can provide on a specific release of the data cannot be considered valid when more data is released. Similarly for sanitised answers to database queries. Despite each query answer may be considered privacy preserving in isolation, this is not the case when considering them together. Intuitively this is due to the correlations between the releases. It should be clear that this problem is present also in LBS offline data release, since the dataset involved in each release is indeed (spatio-temporal) tabular data.

In the case of online data release we can model a threat as a single release when we assume that the adversary can get access only to a single LBS request, or to LBS requests from the same user that cannot be considered correlated (someone assumes this is the case when only sporadic requests are exposed). There are however LBS that require frequent updates of the user location so that their spatio-temporal correlation is clear and can lead to trajectory identification. These cases can be classified as online multiple (or repeated/continuous) data release.

4.1.3 Analysing Privacy Threats

From the definition of privacy threat it is clear how relevant is to understand how the information exchanged as part of the service can be used to identify the individual, to infer private information, and to connect the two to obtain the association that actually leads to the threat.

We have also highlighted the relevance of identifying and modeling the possible adversaries depending on the sharing model, the service architecture, and the type of data release. The prior knowledge that these adversaries may have could be joined with information obtained from the LBS in order to perform a privacy violation. When multiple adversaries may be involved, we should also consider the case of their possible collusion.

As explained later in the chapter, most LBS privacy protection techniques mitigate the risk of privacy threats at an additional cost in terms of computational and communication overhead and/or decrease of service utility (when location information is obfuscated to decrease its sensitivity). Since in LBS service utility is usually dependent on the precision of provided location information, an important parameter to consider when analysing LBS privacy threats is the *precision*, in terms of time and location, required by the service and how the service utility degrades when the precision decreases.

Table 4.1 reports a simplified classification of LBS in terms of required precision, single or continuous release, need of explicit identification and main adversaries (sharing parties). Here we only refer to online data release.

Table 4.1 A classification of location-based services (SP = Service Provider)

LBS type	Req. precision	Continuous	Explicit identity	Adversaries
POI services	High	No	No	SP
Weather/news	Low	No	No	SP
Navigation	High	Yes	No	SP
GeoSN posts	High	No	Yes	SP, users

The service required precision is particularly important to verify the adherence to the principle of *minimization* recommended by many regulators as we will explain in Sect. 4.2.1.

Regarding re-identification, a number of LBS require authentication and profiling, and in these cases we can assume that each request can be straightforwardly associated with an individual, at least by the service provider. In GeoSN, even in case of using pseudonyms, the information existing in public profiles is often sufficient to re-identify individuals. For anonymous services we should carefully consider the re-identification power of the location information transmitted as part of the request. Continuous or repeated requests may also be exploited to re-identify since certain frequent trajectories may be unique to individuals and joined with external information may lead to their identity.

4.1.4 Chapter Organisation

The rest of this chapter is organised as follows. In Sect. 4.2 we discuss regulation compliance and personal preferences as the main requirements for LBS privacy preservation. In Sect. 4.3 we provide an overview of the privacy protection techniques that have been proposed in the literature. We conclude the chapter in Sect. 4.4.

4.2 Compliance with Data Protection Regulation and Individual Privacy Preferences

When designing a LBS or when assessing the privacy impact of an already deployed LBS we need to consider the requirements imposed by the data protection regulation in the countries where the LBS is deployed. Indeed, location privacy protection, similarly to more general privacy protection, is regulated in many countries. In addition to adhering to legal obligations, the design and implementation of LBS should consider also user preferences since different individuals may have very different opinions about sharing their whereabouts with service providers and other users. From a service provider point of view, offering personalised privacy control and transparency through an effective interface may be a competitive advantage.

In the following we briefly introduce these topics providing some references for the interested reader.

4.2.1 A Legal Perspective

The regulation framework for handling geo-location data is fragmented within and across countries, a property unfortunately shared by other types of privacy [30], but particularly true for this type of data. A preliminary analysis of the type of location data being handled should determine if the data is associated with a specific individual and if it is acquired as part of *network traffic* data. In the first case regulations about personal data protection apply, and in the second regulations about data traffic in telecommunication networks also apply. In the following we provide a brief overview of the regulation concerning personal geo-location data, mostly focusing on US and EU.

In the recent past, some specific recommendations for geo-location services have been issued by the EU WP29 Working party, a group of experts and privacy regulation authorities by the EU member states [48]. However, the recommendations did not evolve into a specific regulation but were rather considered in the general data protection regulation. Indeed, in the EU, the principles that should guide the design of a LBS can be extracted from the General Data Protection Regulation (GDPR), approved in May 2016 [47]. Some general principles of the GDPR straightforwardly apply to geo-location data; for example, *privacy by design*, *data portability*, the need of an *informed user consent* (with few exceptions), and the right of the individuals to obtain, update, and even delete their own data. Some other principles can be relatively easily interpreted considering specifically geo-location data; for example, the location privacy interpretation of the *data minimization* principle says that the timestamped geographical position of a user should be acquired and stored only at the precision required for the service being offered. This is well exemplified by localised weather/news services that do not require high precision. Indeed, in a recent recommendation on data processing at work [49], the EU WP29 Working party states that "The information registered from the ongoing monitoring, as well as the information that is shown to the employer, should be minimized as much as possible. Employees should have the possibility to temporarily shut off location tracking, if justified by the circumstances. Solutions that for example track vehicles can be designed to register the position data without presenting it to the employer." Similarly to minimization, *privacy by default* requires the initial location privacy settings of a LBS to be the most protective among the ones available for the specific service. A useful starting point for browsing information about the EU privacy legislation is the *Protection of personal data* web page on the EU official website.[2]

[2] http://ec.europa.eu/justice/data-protection.

In the US, privacy is mostly regulated sector by sector. Regarding location privacy, several U.S. states have enacted laws establishing personal rights. However, current U.S. statute at the federal level does not provide clear protection of geolocation information. The bill known as Geolocation Privacy and Surveillance Act (GPS Act) has been proposed and discussed in congress. As a general principle, the act prohibits companies from collecting or disclosing geolocation information from an electronic communications device without the user's consent. It provides exceptions for parents tracking their children, emergency services, law enforcement, and other cases. Regulations have also been proposed to specifically prohibit development and distribution of "stalking apps," establish an Anti-Stalking Fund at the Department of Justice, and take other steps to prevent geolocation-enabled violence against women. A useful starting point for browsing information about the US is the Geolocation Privacy Legislation page on the GPS.org website.[3]

4.2.2 Privacy Preferences

A successful LBS should not only comply with the applicable regulation, but it should also consider user privacy preferences, and in particular location sharing preferences. Recent experimental studies suggest that the LBS user population has rich location privacy preferences, with a number of critical dimensions, including time of day, day of week, and location [5]. Clearly, another important dimension is who the information is shared with and the precision of the temporal and spatial information being shared. Indeed, the sensitivity of being in a location at a given time is often dependent on the semantics associated with the place, and this can be perceived differently by different individuals.

User preferences can also change over time, not only because users may become more confident in the service and trust more other users or the service provider, but because personal privacy preferences can change based on specific context (e.g., being on a tourist trip with respect to being at home, work or shopping).

The study in [36] also highlights differences between users in different countries in willingness to share location at "home" and at "work" and differences in the granularity of disclosures people feel comfortable with. Several formalisms to represent location privacy policies have been proposed (see e.g., [54]). However, the complexity of control mechanisms offered to LBS users has a clear trade-off with user experience aspects. Complex policy specification interfaces may easily lead to users relying on default settings. Examples of graphical interfaces to set temporal and spatial privacy preferences in mobile location sharing apps can be found in Locaccino [55] and PCube [18]. Some recent efforts have also focused on how to minimize the user intervention in setting and updating privacy preferences in mobile apps using machine learning and other techniques [17, 37].

[3] http://www.gps.gov/policy/legislation/gps-act/.

4.3 Methods and Techniques for Privacy Protection in LBS

In this section, we present and discuss basic methods and techniques for protecting user privacy while using LBS. We restrict ourselves to methods that can be applied *online* since the focus of this chapter is on providing privacy-preserving online LBS, such as services to find points of interest (POI), friend finders in geo-social networks, or online navigation services rather than protecting from the offline analysis of collected data stemming, for instance, from mobility traces collected by a mobile network operator.

Previous surveys on location privacy preserving techniques include [7, 8, 20, 31, 33, 35, 58]. In this section we present a systematized updated view of the research literature in this field.

Protecting user privacy while using LBS implies considering the LBS provider as one of the potential adversaries. Privacy can be protected in two fundamentally different ways:

- **Location-based k-anonymity**. A first approach is to hide the identity of the user since user anonymity would guarantee also the user privacy. Even assuming that we adopt effective anonymization techniques for IP addresses and other general information contained in the LBS queries, the spatio-temporal data contained in the queries can sometimes re-identify the user. In Sect. 4.3.1 we review the research work that has focused on avoiding this re-identification.
- **Sensitive location obfuscation**. Since anonymization is difficult to achieve and sometimes LBS require user identification for using the service, a second possibility to prevent private data leakage is to *restrict the amount of private information being released while interacting with the LBS*. The potentially sensitive information that is specifically released through LBS is the information about the whereabouts of the user: their location at given times. The location may be sensitive by itself (e.g., because the user was supposed to be somewhere else) or may indirectly reveal other sensitive information (e.g., religious or political orientation). In Sect. 4.3.2 we discuss techniques aimed at reducing the sensitivity of the spatio-temporal data in LBS queries by *obfuscating* that information in different ways while trying to preserve the quality of service.

Besides anonymization and obfuscation techniques, we discuss two other classes of approaches: The first uses **cryptographic methods** while the second follows the ideas behind **differential privacy** to devise techniques that provide quantifiable probabilistic guarantees independently from the background knowledge that the adversary may have.

4.3.1 Location-Based k-Anonymity

Anonymity is a general concept not restricted to LBS. So it is a valid question to ask why anonymization for LBS should be different from, say, anonymously visiting

Web pages. The answer is in the fact that location information can be used to de-anonymize users by serving as a so-called *location-based quasi-identifier* [9], which joined with some background knowledge that the adversary may have or acquire, can reveal the user identity. For instance, consider a mobile user called Alice searching for a hospital in her vicinity through a POI finder service. To this end, Alice transmits "anonymously", i.e., without explicitly specifying her true identity in the LBS request, her IP address—as part of the TCP connection—and current location to the LBS to enable the LBS to search for nearby hospitals. In this example, an obvious pseudo-identifier is the IP address of the mobile device of Alice. If the LBS provider has access to a database storing the IP-address-to-customer mapping stored at the Internet service provider of Alice, the LBS can reveal the identity of Alice and might as well conclude that Alice has health problems. Such an attack exploiting the network address can be avoided by using an anonymization service like TOR [13] based on onion routing [24], which forwards the request through a chain of anonymization servers, each one changing the sender IP address for the next "hop" in order not to reveal the initial sender's IP address. However, for an LBS this common anonymization of network addresses is insufficient since the user location might as well serve as quasi-identifier. For instance, assume that Alice sends the request from her home location. Using an easily available map and address book or telephone directory, the LBS can map the home location to Alice's identity, and if Alice is the only person living at that place, this mapping will be unique.

A solution to this problem, which has been applied first by Gruteser and Grunwald to LBS in [25], is the adaptation to the LBS context of the principle of *k-anonymity*. In general, the principle of k-anonymity requires that the individual (here Alice) must be indistinguishable from $k - 1$ other individuals such that the probability of identifying her is $\frac{1}{k}$. Applying k-anonymity in the above example requires the geographic user position to be enlarged to a *cloaking region* including $k - 1$ other users before sending the request to the LBS. This region can be calculated—in the simplest case—by a trusted anonymization service knowing all user positions. Then, even in the worst case in which the untrusted LBS provider can identify all the k users in the reported area, he can only tell that one of them searched for a hospital, and there is only a chance of $\frac{1}{k}$ that this user was Alice.

Obviously, one major challenge of spatial k-anonymity is the calculation of the cloaking region defining the anonymity set containing k users. Intuitively, the size of the area has an impact on the service precision and/or the anonymisation cost. Many different approaches have been proposed to address this problem. Several approaches use a hierarchical spatial partitioning like quadtrees to associate users with cloaking areas [4, 19, 25, 45]. Other approaches use space-filling curves [23] or nearest neighbor (NN) queries with randomization [32] to find groups of k users.

LBS queries generally include timestamps and indeed *location-based quasi-identifiers* are formally defined considering spatio-temporal information [9]. This naturally suggest defenses that generalize time as well as space leading to spatio-temporal cloaking algorithms. An example of temporal cloaking applied in addition to spatial cloaking is the CliqueCloak approach, which proposes to temporarily defer LBS queries [19].

Seeing all these different approaches, the question arises what essential properties are required by a *secure* LBS k-anonymity approach? While we know perfect security is not achievable, we can summarize the important properties that should be fulfilled:

- a cloaking algorithm satisfying reciprocity;
- a mechanism to take correlations among multiple requests into account;
- a trusted anonymizing service or a distributed approach;
- a formal proof that the mechanism is safe with respect to specific assumptions on the adversary's background knowledge.

The last property is the most critical simply because it is difficult to make realistic assumptions on the adversary's background knowledge. This is the main motivation behind the investigation of methods based on the notion of differential privacy that aim at solutions with probabilistic guarantees independent from the adversary's background knowledge. They will be discussed later on, while we now focus on the first three properties.

4.3.1.1 Reciprocity

Considering the first property, it is reasonable to assume that the adversary knows the algorithm for calculating the cloaking region. Then, all the spatial cloaking algorithms should prevent the adversary to use the resulting cloaking region and the algorithm to exclude possible locations of the individual within the region (or equivalently exclude any of the k individuals in the anonymity set). This property has been independently identified and named *reciprocity* by Kalnis et al. [32] and *inversion* by Mascetti et al. [43]. Intuitively, a cloaking algorithm C satisfies this property if each point contained in any cloaking region r computed by C is mapped to r itself.

In the case of k-anonymity, if reciprocity is not fulfilled, the adversary could identify the actual query issuer by executing the algorithm for each of the k users' position and comparing the resulting cloaked region with the one actually received.

4.3.1.2 Correlation Among Multiple Requests

Besides only considering a single query, also *correlation attacks* based on comparing cloaking regions and anonymity sets from *multiple* subsequent queries have been considered. Bettini et al. [9] first identified attacks based on correlating subsequent requests from the same anonymous user and intersecting the corresponding anonymity sets; They introduced the notion of *Historical k-anonymity*. Cloaking algorithms satisfying this property have been designed [43], despite experimental evaluations show that they can deal with very limited temporal sequences of requests after which they need mechanisms to safely change pseudonyms or use other methods to break the correlation.

4.3.1.3 The Diversity Problem

A different form of correlation, based on observing queries from *different* users, also shows the limits of the basic approaches to spatial k-anonymity. The reason follows the motivation for the introduction of l-diversity in data privacy [40]. Bettini et al. [6] illustrated the diversity problem in location k-anonymity by the following example and proposed a way for a trusted anonymiser to compute anonymity sets that avoids this problem.

Consider the LBS user Jane, that is in her office inside a building and she does not want her identity to be associated with the LBS she is requesting. Hence she uses a location k-anonymiser to avoid being re-identified through the coordinates of her own office. Suppose it is known that only Jane and Tom are inside that building at that time. Since she is happy with 2-anonymity, her location is cloaked to the area of the whole building. Suppose that by chance, Tom also asks for the same 2-anonymous service. Since the algorithm satisfies reciprocity the same cloaked region is used in the two requests.Then, the LBS provider, or any adversary that can see the anonymised queries, even if he does not know which of the two requests was sent by Jane will know that Jane (and Tom) asked for that (sensitive) service.

A slightly different notion of location l-diversity is introduced by Bamba et al. [4] for the case in which the location information in the LBS queries is not re-identifying but it is *sensitive* information.

Consider the LBS user Bob searching for a nearby taxi through a POI finder. Assume that Bob is currently located at a hospital. Even if there are $k-1$ other patients at the hospital, a LBS query from Bob's current location (the hospital) would reveal that Bob might have health problems. The problem here is that all k users are located at the same sensitive location (hospital).

To avoid such inference, their PrivacyGrid approach applies spatial cloaking so that l different symbolic addresses (e.g., hospitals and *other* types of locations) are within the same cloaking region.

4.3.1.4 Trusted Anonymiser or Distributed Anonymization

Using a centralized service to calculate a cloaking region containing k potential users requires this service to be trusted, and introduces a potential bottleneck and single-point of failure with respect to availability and—more importantly—privacy. In particular, a compromised anonymization service reveals all true user positions. Therefore, decentralized approaches for spatial k-anonymity based on Hilbert space-filling curves have been proposed by Ghinita et al. [22, 23], to calculate anonymity sets distributedly between a set of "peers". Further distributed peer-to-peer approaches have been proposed, e.g., based on measuring the distance to other peers using WiFi signal strength and a scheme for distributively calculating the cloaking region by peers without revealing precise location information to other peers [29].

Among the proposals that do not require a trusted anonymiser, Kido et al. [34] proposed to locally generate, for each LBS query, a set of position dummies (*fake* locations) as the locations of other users and sending them to the LBS together with the true location. This is equivalent to issue multiple LBS queries for each original query. The results form the service are then locally filtered to retain only the ones related to the correct user location. The intention of the authors is to achieve a form of k-anonymity using $k - 1$ dummies, but the limited local knowledge does not guarantee that the generated fake positions are actual position of other potential users. Moreover, the adversary may use public information about geographical and street network constraints to exclude some of the dummy positions. The SybilQuery technique by Shankar et al. [51] follows the same approach, but it improves the quality of the dummies by using locations with similar traffic conditions, exploiting a database containing historic traffic, and traffic restrictions like one-way streets. However, the technique has similar limits.

4.3.2 Protecting Location Information by Obfuscation and Perturbation

In contrast to the anonymization techniques discussed in the previous section, obfuscation techniques do not try to hide the identity of the user. Indeed, there are several LBS requiring authentication and others for which the identity may be easily derived from other information in the request. Instead, these techniques aim at blurring or perturbing the location information contained in LBS requests because of its potential sensitivity.

As an example, the precision of location information can be decreased by translating precise point coordinates to geographic regions; Analogously, the precision of the temporal information usually associated with location can be decreased by converting precise timestamps into time intervals. Note that, as opposed to anonymity approaches, in this case the location is not enlarged in order to include other potential users, but to decrease the sensitivity, for example by including different types of semantic locations. This is a fundamental difference between the two protection approaches.

Although most approaches reduce precision by using areas that contain the user location, some approaches also reduce accuracy by sending a fake location, which might be specified very precisely, e.g., as a point coordinate, but deviates from the true user location.

Sending to a LBS inaccurate location and time information may impact the quality of service (QoS) of the LBS. For instance, searching for the nearest restaurant through an LBS might yield different (imprecise) results when given a larger obfuscation area rather than the precise position of the user. However, the QoS of the final result can be improved by post-processing imprecise or inaccurate

results returned by the LBS by the same agent that performed the obfuscation, as for example, the app on the user's smartphone.

As a positive property compared to anonymization, obfuscation typically does not require an additional trusted infrastructure like, for instance, a trusted anonymization service as used by the centralized k-anonymity approaches discussed above. Instead, obfuscation (and answer filtering) can be performed on the user device alone, possibly assisted by some locally available information like maps.

Similarly to anonymity, an obfuscation defense should satisfy reciprocity, i.e., we must assume that the adversary knows the obfuscation algorithm, and it should be proved that this knowledge would not lead to any privacy leak. For instance, assume a naive deterministic obfuscation algorithm that simply creates a circular area centered at the true user position. Obviously, in this example an adversary can simply revert the obfuscation since the distribution of the user position is very skewed with a 100% chance of the user being at the center of the circle. Therefore, one of the challenges in designing a defense is to devise an algorithm withstanding attacks calculating the probability density function of the user location.

Jensen et al. [31] provide a good survey of obfuscation based techniques. In the following we summarise these techniques by grouping them as follows:

- Query enlargement techniques
- Dummy based techniques
- Coordinate transformation techniques

4.3.2.1 Query Enlargement Techniques

We group in this class all techniques that instead of including a specific location (and time) as part of the LBS query, they include a geographical area (and a time instant/interval) often called *obfuscation area*.

Ideally, the distribution of the user position within the spatio-temporal obfuscation area should be uniform not to give the adversary any hint about where the user might be located. However, this is not trivial to achieve since spatial constraints like streets, buildings, lakes, or forests might increase or decrease the probability of users being located in some parts of the obfuscation area. Correlations with other queries could also rule out some spatio-temporal areas.

A representative obfuscation approach based on transforming user positions in circular areas by applying a set of operators is the one proposed by Ardagna et al. [3]. The obfuscation operators can enlarge the circle, shifting the center of the circle, and shrinking the circle; the effects of applying these operators on the probability distribution of the user position is analyzed. Randomness can be introduced, for instance, by shifting the circle into a random direction. That work has been also extended by the authors considering background knowledge such as maps that might assist an adversary to find locations within the obfuscation area where the user might be located with higher probability. The obfuscation area is adapted, for instance, by

increasing the radius of the circle, to compensate for the constraints given by a map leading to a non-uniform distribution over the obfuscation area.

Damiani et al. [12] also considered background knowledge in their obfuscation approach modeling the fact that positions are not uniformly distributed. Beyond this, they also considered the semantics of locations for calculating obfuscation areas, leading to different sensitivity of locations from a user's perspective. For instance, a user might not want to disclose that he is currently in a hospital, thus, the obfuscation area should include other non-sensitive locations leading to a low probability of the user being located in a hospital.

Spatial k-Group-Nearest-Neighbor (kGNN) queries over obfuscated location information are in the focus of [27]. Such a kGNN returns the "meeting point" minimizing the aggregated distance to all group members. For instance, an LBS could propose a restaurant minimizing the travel distance of a geographically distributed group of people. The privacy objective here is not to reveal the precise location to the LBS nor to other group members. To this end, each group member obfuscates his location by a rectangular area, which is sent to the LBS. Given imprecise locations of all group members, the LBS can only calculate a candidate set for the kGNN. This candidate set is post-processed by each group member sequentially to calculate the final kGNN.

A computational method that can be used both for anonymization and obfuscation through query enlargement has been presented by Mascetti et al. [42]. The method is agnostic about the semantics of the generalization function (for anonymity the semantics concerns the number of candidate individuals in the region, for obfuscation it may be the size of the area, the type of the area, the number of different pubs, etc...). Moreover, as opposed to most reciprocity-safe methods for finding generalized regions, it does not partition the space but uses an efficient bottom-up approach to find for each LBS query its generalised spatial region, called *Safebox*.

The trade-off between QoS and privacy as achieved by obfuscation is studied by Cheng et al. in [1]. Obviously, the definition of QoS essentially depends on the service offered by the LBS. In this work, the authors consider spatial range queries as a primitive frequently used by LBS. The authors assume that both the location of the query issuer as well as the locations of the queried objects, are obfuscated. In that case, answers to the range query are probabilistic, since some obfuscated locations might overlap with the queried range, and a precise answer by the LBS about the object being inside or outside the queried range is not possible.

4.3.2.2 Dummy Based Techniques

The idea of generating fake user positions proposed by Kido et al. [34] for k-anonymity can also be used to hide a possibly sensitive user location. The generation algorithm in this case has a different goal: the fake position should not resemble the position where another potential user is located, but it should be a non-sensitive location where the actual user could be. For example, instead of reporting a medical

facility, the address of a grocery store across the street (or a set of such locations) is reported as the current location in the LBS query. Several dummy generation algorithms in this category are proposed by Lu et al. [39].

Related to dummy locations is the approach proposed by Duckham and Kulnik [14]; They show how to apply obfuscation to graph models. Graph vertices model locations, including the current user location, while edges model connections between locations such as roads, which can also have a weight to model some notion of distance. In their approach, the current location is obfuscated by sending as part of the LBS query a set of vertices representing dummy locations plus the actual location. Clearly, the more elements are contained in the set, the more imprecise is the obfuscated location. The LBS answers to proximity queries asking for the closest resource by performing computations on the graph model. The authors propose a negotiation protocol by which the LBS can ask for a smaller set of candidate locations in order to improve the quality of service.

Finally, the SpaceTwist approach by Yiu et al. [61] addresses the location privacy problem in answering k-Nearest-Neighbor (kNN) queries with a dummy-based *progressive retrieval technique*. Indeed, it generates a single dummy location called *anchor* and communicates only that location to the LBS. The distance of the dummy location from the real one is a parameter and it determines the achieved level of privacy. The SpaceTwist algorithm incrementally queries the LBS about the nearest objects for the same given anchor. These results are then filtered on client-side to find the actual k nearest neighbors for the true user position.

4.3.2.3 Coordinate Transformation Techniques

Another obfuscation method that has been explored is *coordinate transformations* [26]. Instead of creating obfuscation areas or dummy locations, coordinate transformations change the complete coordinate reference system using geometric transformations such that transformed coordinates cannot be interpreted by the adversary with respect to a "real-world" location on earth. However, the transformation should still allow for the LBS to answer the queries. For example, a friend finder service should still be able to evaluate proximity, i.e., the transformation should, at least approximately, preserve the distance. For instance, in [26] the authors outline how to use coordinate transformations for implementing basic spatial queries such as position queries, spatial range queries, and to detect spatial events such as "on entering area" or "on meeting" events. The idea is that the LBS managing mobile user positions performs query processing on transformed coordinates, while the transformation rules serve as shared secret between a user and other users or services with whom the user wants to share his location.

The essential challenge for coordinate transformation approaches is that an adversary can exploit background knowledge like maps and spatial distributions of locations to revert the transformation, i.e., to find the original location on earth given the transformed coordinates. In [38], the authors analyze distance preserving transformations as proposed in privacy preserving data mining. They conclude

that approximate locations of users can be inferred based only on partial relative distance information and publicly available background knowledge about mobile object distributions. A specific attack to LBS protected by distance preserving transformations has been shown to be practical by Mascetti et al. [41].

4.3.3 PIR and Cryptographic Approaches

Private information retrieval (PIR) and cryptographic methods, namely, encryption, cryptographic hashing, secret sharing, and secure multi-party computation have also been considered to implement privacy-aware LBS. The basic objectives of these approaches are the same as for the approaches discussed above, namely, anonymity and sensitive location protection. However, by applying proven cryptographic methods, these approaches strive for stronger, provable privacy guarantees. The essential challenge is to allow for efficient processing of spatial queries at the LBS provider, although location information is not available in plain text to the provider (encryption and hashing methods), or despite the computational complexity of the cryptographic method (private information retrieval, secure multi-party computation).

Ghinita et al. in [21]. apply the concept of *Private Information Retrieval* (PIR) to an LBS implementing spatial nearest neighbor (NN) queries. The general idea of PIR is to privately retrieve data from a database without revealing which information has actually been requested. Applied to spatial NN queries, the goal is to retrieve the objects (POIs) nearest to the query issuer without revealing to the LBS which spatial region has actually been queried by the user. A naive solution would be to query the whole database, i.e., all POIs, however, obviously the overhead would be very high. Informally, PIR reduces this overhead by sending an encrypted query to the LBS not revealing what entry has been queried, but allowing the LBS to return a result significantly smaller than the whole database to the client, which then can be used by the client to calculate the value of the actually queried database entry. For mapping POIs to database entries, the authors use space-filling curves to preserve the spatial proximity required by NN queries. Khoshgozaran and Shahabi [33] provide a comprehensive survey of PIR approaches to LBS privacy preservation. Despite the solid theory, the PIR techniques have not yet been proven practical and scalable mainly for efficiency reasons.

In the application area of social network LBS, and in particular friend-finders, *secure multi-party computation* (SMC) has been used to implement protocols for computing proximity [62]. More generally, the basic objective of SMC is to jointly calculate a known function (e.g., proximity) by n participants, each participant providing a secret input to the function (e.g., position), without revealing the secret input to the other participants.

A cryptographic approach targeting location privacy in friend-finder services has been proposed by Mascetti et al. in [44]. The objective of this approach is to allow participants to issue queries to a central service for finding all friends

within a given distance, while hiding to the service provider any information about their position and proximity of other users. Their method also allow the user to control the precision of the location information released to friends. To this end, proximity is computed by using a combination of cryptographic hash functions and SMC exploiting the commutative property of an encryption function. Location information is encrypted at different levels of granularity so that, according to privacy preferences, friends will only be able to infer the user's position with a given approximation. The system has been implemented in a prototype app, called *PCube*, that has been available both for iOS and Android devices.

Another cryptographic approach is based on the concept of *secret sharing*. The basic idea of secret sharing is to split a secret into a number of shares, say n. The secret can be revealed if a certain number of shares, say t, are known (so-called (t, n)-threshold scheme [50]). This concept is applied in [57] to implement a distributed location service managing locations of a user population and providing a set of LBSs with location information. To this end, locations—which can be geographic or symbolic locations—are defined as secrets. n shares are generated per location and distributed among n different servers. Consequently, in order to reveal the location, an adversary has to break into t servers, thus, avoiding a single point of failure. Moreover, by using a *multi-secret sharing scheme*, this approach supports providing location information of different precision levels, corresponding to a multi-secret, to different LBSs querying the location service.

4.3.4 Differential Privacy Approaches

Considering the difficulty of providing formal privacy guarantees independent from background knowledge for anonymity and obfuscation based approaches as well as the costs and applicability limitations of cryptographic approaches, a new type of methods has been proposed inspired by the success of the *differential privacy* notion in statistical databases.

Differential privacy has been introduced by Dwork [15] in statistical databases as a general method for the privacy preserving analysis of tabular personal data. The intuitive idea behind differential privacy is the following: Given two databases that differ only for the second including an additional record about an individual that is not present in the first, the information separately extracted from the two databases with a differential privacy method will not be significantly different. In other words the result of the analysis will be independent from the presence of information about the specific individual, hence it cannot be used in any way to violate her privacy. The way this result is achieved is by probabilistically inserting noise in the data. We refer to the original paper and to the rich literature on this topic, including other contributions in this book, for a formal definition and technical properties.

Differential privacy had a significant impact also on location privacy with natural applications to the offline analysis of location data, as in answering counting queries on a large dataset of user positions [11]. He et al. have also shown how

to use differential privacy methods to synthesize mobility data based on raw GPS trajectories of individuals while ensuring strong privacy protection [28].

4.3.4.1 Differentially Private Methods for LBS

A more challenging task has been adapting the principles of differential privacy to online data release in LBS. (D, ϵ)-location privacy, illustrated by Elsalamouny and Gambs [16] results from adapting the adjacency relation in the standard differential privacy to the domain of locations. Two locations are considered "adjacent" if the distance between them is less than a predefined value D. In this context, a mechanism satisfies(D, ϵ)-location privacy if the (log of) the ratio between the probabilities of obtaining a certain output, from any two adjacent locations is at most ϵ. This property guarantees that the distinguishability between the location of the user and all the points that are adjacent is always restricted to a certain level quantified by ϵ.

A similar extension of differential privacy introduced by Andres et al. [2] is ϵ-geo-indistinguishability in which the bound on the distinguishability between two arbitrary positions increases linearly with the distance d between them. This means that the (logarithm of) the ratio between the probabilities of obtaining a certain output from two locations is at most d, which provides a low level of distinguishability (i.e., high privacy) between neighboring positions. In contrast, a higher level of distinguishability (i.e., low privacy) occurs for points that are further apart.

Analogously to the original differential privacy proposal, the way to achieve these properties is by inserting noise. In the LBS case this is done by *probabilistically determining a fake location that replaces the real location* when performing the LBS query. Different randomization functions can be used as long as they allow to prove the desired differential properties [16]. In the original proposal of ϵ-geo-indistinguishability *planar Laplacian* noise is used. The investigation of alternative randomization functions with more favorable trade-offs between privacy and utility is an active research area.

Finally, analogously to what we have seen for anonymity and obfuscation, differential privacy methods also have to deal with multiple (sequential) release of data and, more generally, with correlations that an adversary may exploit. A composition theorem for differential privacy says that we should consider the sum of the ϵ values associated with each release in the sequence. When considering LBS that require frequent or continuous queries this seems to imply that we would quickly reach unacceptable values of ϵ. A result consistent with what has been experimentally observed with the spatial cloaking for anonymity. An Interesting work on protecting locations from temporal correlations under differential privacy has been done by Xiao et al. [59, 60].

4.3.4.2 Analysing Trade-Offs Between Protection and Utility

The major critic to differential privacy is on the practical utility of the resulting mechanisms since keeping the ϵ parameter low happens at the expenses of the utility of the resulting query answers, which in the domain of LBS is the quality of service. A number of research efforts are directed to investigate this problem [10]. A natural question that still has no clear answer is what value of ϵ provides a good level of privacy. More theoretical work considers as good values close to 1, while applications seem to use quite higher values.

As part of research on finding optimal trade-offs between privacy and utility, Shokri [52] proposes a game theoretic approach to find an optimal location protecting mechanism while respecting each individual user's service quality requirements. Protection is achieved by a combination of differential privacy and distortion functions.

The application of game theory cited above for finding optimal trade-offs between privacy and quality of service has also been extended to deal with multiple releases, i.e., sets of queries that may reveal location traces [53].

4.4 Conclusions

In this chapter we provided a classification of the many location based services that are being offered today, and we illustrated the privacy threats that their users may face when using these services.

An important message that we would like to convey is that in order to understand if a given privacy protection method is adequate for a given service, it is necessary to carefully analyse the service in terms of the information being exchanged, the service architecture and the different parties to which the information is exposed, as well as to evaluate the requirements in terms of location data accuracy in order not to degrade the service quality. We also highlighted the importance of modeling the adversaries in terms of their access to LBS queries (single or multiple queries, sporadic or continuous) and in terms of the prior knowledge that they may have or acquire, including the knowledge of the privacy preserving algorithm and parameters.

In this chapter we also briefly reviewed the legal framework and personal privacy preferences as hard and soft requirements to be considered in the design of a defense technique. Finally, we provided a survey of the technical solutions proposed for on-line protection of LBS queries, hence focusing on techniques that aim to protect personal data before they reach the service provider, as opposed to offline techniques that aim at protecting the release of datasets from LBS providers to third parties, typically for statistical analysis. Other chapters in this book provide a deeper coverage of some of the approaches illustrated in Sect. 4.3 when applied to specific categories of LBS.

Overall, we can conclude observing that the online location privacy protection problem is a very challenging one, especially if considering the protection of trajectories as revealed by sequences of correlated LBS queries. The difficulty is mostly due to the uniqueness property of human trajectories [46] and to modeling realistic assumptions about the prior knowledge of the adversary.

Despite the protection proposals based on the notion of differential privacy have the advantage of providing provable probabilistic guarantees independent from the adversary's knowledge, their utility in terms of quality of service for many LBS is still to be demonstrated. An interesting research direction would be considering new probabilistic methods to insert location noise based on specific LBS deployment contexts, user preferences, and adversary model. Some of the research results obtained by the anonymization and obfuscation approaches may turn out to be applicable.

Acknowledgements We would like to thank Frank Dürr for providing a preliminary draft of the description of some of the established defense techniques as analysed in his survey [58].

References

1. Preserving user location privacy in mobile data management infrastructures. *Privacy Enhancing Technologies*, LNCS 4258:393–412, 2006.
2. M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. Geo-Indistinguishability: Differential Privacy for Location-Based Systems. In *Proceedings of ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2013.
3. C. A. Ardagna, M. Cremonini, S. D. Capitani, and P. Samarati. An Obfuscation-based Approach for Protecting Location Privacy. *IEEE Transactions on Dependable and Secure Computing (TDSC)*, 8(1):13–27, 2011.
4. B. Bamba, L. Liu, P. Pesti, and T. Wang. Supporting anonymous location queries in mobile environments with PrivacyGrid. In *Proc. Int. WWW Conf.*, pages 237–246, 2008.
5. M. Benisch, P. G. Kelley, N. Sadeh, and L. F. Cranor. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing*, 15(7):679–694, 2011.
6. C. Bettini, S. Jajodia, and L. Pareschi. Anonymity and diversity in LBS: A preliminary investigation. In *IEEE International Conference on Pervasive Computing and Communications Workshops*, 2007.
7. C. Bettini, S. Mascetti, D. Freni, X. S. Wang, and S. Jajodia. Privacy and anonymity in Location Data Management. In F. Bonchi and E. Ferrari, editors, *Privacy-Aware Knowledge Discovery: Novel Applications and New Techniques*. Chapman & Hall, 2010.
8. C. Bettini, S. Mascetti, X. Wang, D. Freni, and S. Jajodia. Anonymity and historical-anonymity in location-based services. In *Privacy in location-based applications*, volume LNCS 5599. Springer Berlin Heidelberg, 2009.
9. C. Bettini, X. Wang, and S. Jajodia. Protecting privacy against location-based personal identification. In *Proceedings of Secure Data Management*, volume 3674 LNCS. Springer, 2005.
10. K. Chatzikokolakis, E. Elsalamouny, and C. Palamidessi. Efficient Utility Improvement for Location Privacy. *Proceedings on Privacy Enhancing Technologies (PoPET)*, 2017(4):210–231, 2017.

11. G. Cormode, C. Procopiuc, D. Srivastava, E. Shen, and T. Yu. Differentially private spatial decompositions. In *Proceedings - International Conference on Data Engineering*, pages 20–31, 2012.
12. M. L. Damiani, E. Bertino, and C. Silvestri. The PROBE framework for the personalized cloaking of private locations. *Transactions on Data Privacy*, 3(2):123–148, 2010.
13. R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. Technical report, Naval Research Lab Washington DC, 2004.
14. M. Duckham and L. Kulik. A Formal Model of Obfuscation and Negotiation for Location Privacy. In *International Conference on Pervasive Computing*, pages 152–170. Springer Berlin Heidelberg, 2005.
15. C. Dwork. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008.
16. E. Elsalamouny and S. Gambs. Differential Privacy Models for Location- Based Services. *Transactions on Data Privacy*, 9:15–48, 2016.
17. K. Fawaz and K. G. Shin. Location privacy protection for smartphone users. In *ACM Conference on Computer and Communications Security*, 2014.
18. D. Freni, S. Mascetti, C. Bettini, and M. Cozzi. Pcube: A system to evaluate and test privacy-preserving proximity services. In *2010 Eleventh International Conference on Mobile Data Management*, pages 273–275, May 2010.
19. B. Gedik and Ling Liu. Location Privacy in Mobile Systems: A Personalized Anonymization Model. In *25th IEEE International Conference on Distributed Computing Systems (ICDCS'05)*, pages 620–629, 2005.
20. G. Ghinita. Privacy for location-based services. *Synthesis Lectures on Information Security, Privacy & Trust*, 4(1):1–85, 2013.
21. G. Ghinita, P. Kalnis, A. Khoshgozaran, C. Shahabi, and K.-L. Tan. Private queries in location based services: anonymizers are not necessary. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 121–132. ACM, 2008.
22. G. Ghinita, P. Kalnis, and S. Skiadopoulos. MOBIHIDE: a mobilea peer-to-peer system for anonymous location-based queries. *Advances in spatial and temporal databases*, pages 221–238, 2007.
23. G. Ghinita, P. Kalnis, and S. Skiadopoulos. PRIVE: anonymous location-based queries in distributed mobile systems. In *Proceedings of the 16th international conference on World Wide Web*, pages 371–380. ACM, 2007.
24. D. Goldschlag, M. Reed, and P. Syverson. Onion routing. *Communications of the ACM*, 42(2):39–41, 1999.
25. M. Gruteser and D. Grunwald. Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. In *Proceedings of the 1st international conference on Mobile systems, applications and services - MobiSys '03*, pages 31–42, 2003.
26. A. Gutscher. Coordinate transformation-a solution for the privacy problem of location based services? In *Parallel and Distributed Processing Symposium, 2006. IPDPS 2006. 20th International*, pages 7—-pp. IEEE, 2006.
27. T. Hashem, L. Kulik, and R. Zhang. Privacy preserving group nearest neighbor queries. In *Proceedings of the 13th International Conference on Extending Database Technology*, pages 489–500. ACM, 2010.
28. X. He, G. Cormode, A. Machanavajjhala, C. M. Procopiuc, and D. Srivastava. DPT: Differentially Private Trajectory Synthesis Using Hierarchical Reference Systems. *Proceedings of the VLDB Endowment*, 8(11):1154–1165, 2015.
29. H. Hu and J. Xu. Non-exposure location anonymity. In *Proceedings - International Conference on Data Engineering*, pages 1120–1131, 2009.
30. R. P. Jay. Data protection & privacy in 31 jurisdictions worldwide. Gideon Roberton, Law Business Research Ltd, 2015.
31. C. S. Jensen, H. Lu, and M. L. Yiu. Location privacy techniques in client-server architectures. In *Privacy in location-based applications*, pages 31–58. Springer, 2009.

32. P. Kalnis, G. Ghinita, K. Mouratidis, and D. Papadias. Preventing location-based identity inference in anonymous spatial queries. *IEEE Transactions on Knowledge and Data Engineering*, 19(12):1719–1733, 2007.
33. A. Khoshgozaran and C. Shahabi. Private information retrieval techniques for enabling location privacy in location-based services. In *Privacy in Location-Based Applications*, volume 0831505, pages 59–83. Springer Berlin Heidelberg, 2009.
34. H. Kido, Y. Yanagisawa, and T. Satoh. Protection of location privacy using dummies for location-based services. In *21st International Conference on Data Engineering Workshops.*, page 1248. IEEE, 2005.
35. J. Krumm. A survey of computational location privacy. *Personal and Ubiquitous Computing*, 13(6):391–399, 2009.
36. J. Lin, M. Benisch, N. Sadeh, J. Niu, J. Hong, B. Lu, and S. Guo. A comparative study of location-sharing privacy preferences in the United States and China. *Personal and Ubiquitous Computing*, 17(4):697–711, Apr 2013.
37. B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. Zhang, N. Sadeh, A. Acquisti, and Y. Agarwal. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Symposium on Usable Privacy and Security*.
38. K. Liu, C. Giannella, and H. Kargupta. An attacker's view of distance preserving maps for privacy preserving data mining. *Knowledge Discovery in Databases: PKDD 2006*, pages 297–308, 2006.
39. H. Lu, C. S. Jensen, and M. L. Yiu. PAD: Privacy-Area Aware, Dummy-Based Location Privacy in Mobile Services. In *Proceedings of the CM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDE)*, pages 16–23, 2008.
40. A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkitasubramaniam. l-diversity : Privacy beyond k-anonymity. In *Proceedings - International Conference on Data Engineering*, 2006.
41. S. Mascetti, L. Bertolaja, and C. Bettini. A practical location privacy attack in proximity services. In *Proceedings - IEEE International Conference on Mobile Data Management*, 2013.
42. S. Mascetti, L. Bertolaja, and C. Bettini. SafeBox : adaptable spatio-temporal generalization for location privacy protection. *Transactions on Data Privacy*, 7:131–163, 2014.
43. S. Mascetti, C. Bettini, D. Freni, and X. S. Wang. Spatial generalisation algorithms for LBS privacy preservation. *Journal of Location Based Services*, 1(3):179–207, 2007.
44. S. Mascetti, D. Freni, C. Bettini, X. Wang, and S. Jajodia. Privacy in geo-social networks: Proximity notification with untrusted service providers and curious buddies. *VLDB Journal*, 20(4), 2011.
45. M. F. Mokbel, C.-Y. Chow, and W. G. Aref. The new casper: Query processing for location services without compromising privacy. In *Proceedings of the 32nd international conference on Very large data bases*, pages 763–774. VLDB Endowment, 2006.
46. Y.-a. D. Montjoye, M. Verleysen, and V. D. Blondel. Unique in the Crowd: The privacy bounds of human mobility. *Scientific Reports*, 3(1376):1–5, 2013.
47. E. Parliament and Council. General data protection regulation - 2016/679. Technical report, European Commission, 2016.
48. A. D. P. W. Party. Opinion 13/2011 on geolocation services on smart mobile devices. Technical report, European Commission, 2011.
49. E. A. D. P. W. Party. Opinion 2/2017 on data processing at work - wp249. Technical report, European Commission, 2017.
50. A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
51. P. Shankar, V. Ganapathy, and L. Iftode. Privately querying location-based services with SybilQuery. In *Proceedings of the 11th international conference on Ubiquitous computing*, pages 31–40. ACM, 2009.
52. R. Shokri. Privacy Games : Optimal User-Centric Data Obfuscation. In *Proceedings on Privacy Enhancing Technologies*, pages 299–315, 2015.
53. R. Shokri, G. Theodorakopoulos, and C. Troncoso. *ACM Transactions on Privacy and Security (TOPS)*.

54. E. Snekkenes. Concepts for personal location privacy policies. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 48–57. ACM, 2001.
55. E. Toch, J. Cranshaw, P. Hankes-Drielsma, J. Springfield, P. G. Kelley, L. Cranor, J. Hong, and N. Sadeh. Locaccino: A privacy-centric location sharing application. In *Proceedings of the 12th ACM International Conference Adjunct Papers on Ubiquitous Computing - Adjunct*, UbiComp '10 Adjunct, pages 381–382, New York, NY, USA, 2010. ACM.
56. C. R. Vicente, D. Freni, C. Bettini, and C. S. Jensen. Location-related privacy in geo-social networks. *IEEE Internet Computing*, 15(3):20–27, 2011.
57. M. Wernke, F. Durr, and K. Rothermel. PShare: Position sharing for location privacy based on multi-secret sharing. In *2012 IEEE International Conference on Pervasive Computing and Communications*, pages 153–161, 2012.
58. M. Wernke, P. Skvortsov, F. Durr, and K. Rothermel. A classification of location privacy attacks and approaches. *Personal and Ubiquitous Computing*, 18(1):163–175, 2014.
59. Y. Xiao and L. Xiong. Protecting Locations with Differential Privacy under Temporal Correlations. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1298–1309, 2015.
60. Y. Xiao, L. Xiong, S. Zhang, and Y. Cao. LocLok: Location Cloaking with Differential Privacy via Hidden Markov Model. *Proceedings of the VLDB Endowment (VLDB)*, 10(12):1901–1904, 2017.
61. M. L. Yiu, C. S. Jensen, X. Huang, and H. Lu. SpaceTwist: Managing the Trade-Offs Among Location Privacy, Query Performance, and Query Accuracy in Mobile Services. In *IEEE International Conference on Data Engineering (ICDE)*, pages 366–375, 2008.
62. G. Zhong, I. Goldberg, and U. Hengartner. Louis, lester and pierre: Three protocols for location privacy. In *Privacy Enhancing Technologies*, pages 62–76. Springer, 2007.