Chapter 14
Location Privacy-Preserving Applications and Services

Ioannis Boutsis and Vana Kalogeraki

Abstract Mobile location-based applications have recently prevailed due to the massive growth of the mobile devices and the mobile network. Such applications give the opportunity to the users to share content with the community which is coupled with their current geographical location. However, sharing such information might have serious privacy implications as an adversary might monitor the system and use such information to expose sensitive user information including user mobility traces and sensitive locations. This problem has led both the research community and the commercial mobile applications to develop several solutions to handle these privacy implications so as to enable users to disclose content without compromising their privacy. This chapter provides a survey of the state-of-the-art location-based mobile applications, describes the privacy implications that arise from contributing information in such applications and the respective existing countermeasures to deal with the privacy preservation issues. Furthermore, we describe our experiences from deploying a real-world location-based application that aims to allow the user contribute content and protect the user's privacy.

14.1 Introduction

Mobile applications have recently become a core part of commercial systems targeted to end-users, as they enable users to access their services from everywhere. Mobile applications typically provide the same features with the respective desktop applications, but they are developed for mobile devices such as smartphones and tablets. However, the vast majority of the mobile applications also take advantage of the ability to acquire the user location from the mobile devices' sensors, a feature which is not available in traditional desktop applications, in order to generate

I. Boutsis · V. Kalogeraki (✉)
Athens University of Economics and Business, Athens, Greece
e-mail: mpoutsis@aueb.gr; vana@aueb.gr

A. Gkoulalas-Divanis, C. Bettini (eds.), *Handbook of Mobile Data Privacy*,
https://doi.org/10.1007/978-3-319-98161-1_14

personalized content, depending on the respective location, and provide a richer experience compared to the desktop version of the service.

From the user perspective, it seems that more and more users appreciate mobile applications that exploit the geo-location features and are willing to share their location information with the service provider and potentially with other users, despite their privacy concerns. This is due to the fact that location-based applications provide additional features as they enable them to share their location with their acquaintances, including their friends and family, and to explore content which is related to their current location, such as weather, news, nearby places or nearby friends.

Although users are tempted to share their location information to enjoy the location-based services, sharing such data makes them vulnerable to several types of privacy attacks. This is because location data constitutes sensitive user information and the nature of the location information enables an adversary to expose a great amount of information per user with only a small set of geo-located information. Hence, if this data is exploited from potentially untrusted parties it can lead to severe consequences that range from user profiling for advertising purposes to real-world crimes such as stalking, robberies, etc.

This fact has led both the research community as well as the commercial mobile applications to develop several privacy preservation approaches to cope with this problem. However, as we explain in this chapter the commercial mobile applications typically use simple approaches which are prone to attacks rather than exploiting the complex but effective approaches that have been proposed in the research literature.

The goal of this chapter is to provide a survey of the most popular location-based applications and state-of-the-art privacy mechanisms. The rest of this chapter is organized as follows. Section 14.2 presents the most popular location-based applications and explains the benefits of sharing location-based data. Section 14.3 describes the characteristics of the shared data and the respective privacy implications that derive from these characteristics. Section 14.4 summarizes the existing privacy mechanisms that have been developed as countermeasures. Section 14.5 gives a discussion of our experiences from employing a real-world location based application and the privacy mechanism that we used. Finally, Sect. 14.6 concludes the chapter.

14.2 Popular Location Sharing Applications

The first section of this chapter introduces the most popular location sharing applications. In this chapter we focus on the following categories:

1. Social Applications
2. Transportation Applications
3. Travel Applications

4. Fitness Applications
5. Image Sharing Applications
6. Location Sharing Applications initiated by the research community and focus on preserving the user privacy

All the location sharing applications that are presented in this chapter have been selected due to their popularity and constitute the most prominent location sharing applications at the time of writing this chapter.

14.2.1 Social Applications

Social services have recently become extremely popular as they enable users to connect and interact with their friends and family. Social services give the opportunity to the users to create social profiles and share content such as text posts and photos or videos with their social connections. They also facilitate the development of online social networks by connecting a user's profile with those of other individuals and/or groups.

The prevalence of mobile technologies has enabled the social services to develop mobile apps that exploit the embedded sensors in order to provide a richer experience for the users. Hence, they allow the users to include location-embedded information or content. Introducing location capabilities, such as geotagging, provides additional features to the users such as letting their acquaintances know where they are or where a specific image was captured. However, revealing the physical location of the user can lead to significant privacy implications.

14.2.1.1 Facebook

Facebook (http://www.facebook.com) is one of the prominent location-based social networks with 1.66 billion mobile monthly active users (as of September 30, 2016).[1] It allows users to share various types of content with their designated set of users, including posts, photos and videos and to chat through Messenger with the rest of the Facebook users.

Each type of content which is shared by a user in Facebook from the mobile app may be coupled with the user geo-location. Hence, except from social posts, the user can also share her geo-location even during a chat. However, Facebook includes a consistent indicator as a reminder when the users share their location.

In Facebook, the content published by the users can be visible by all the users who are authenticated to access the shared content. Thus, the user has the

[1] http://newsroom.fb.com/company-info/.

responsibility to determine her audience and manage her privacy by assigning the users and groups that should be available to view shared content.

14.2.1.2 Twitter

Twitter (http://www.twitter.com) is another well-known social service used by 313 million monthly active users, from which 82% are mobile.[2] The idea behind Twitter is that it allows users to send short 140-character messages called "tweets" to interact with the community. Registered users can post and read tweets, but those who are unregistered can only read them. Mobile users are also able to embed their geographical location in these tweets to share the location where the tweet was produced.

Twitter provides two levels of privacy for the tweets that can be selected by the users: (1) Public Tweets and (2) Protected Tweets. Public Tweets, which is the default setting, makes the tweets visible to anyone, even users that do not own a Twitter account. On the other hand, Protected Tweets can only be visible by users that have been authenticated from the producer of the tweet so as to protect user privacy.

14.2.1.3 Foursquare

Foursquare (http://www.foursquare.com) is a leading location-based social networking website for mobile devices with more than 50 million people using Foursquare each month, through the web service and the mobile app.[3] Foursquare allows registered users to post their attendance at a venue (referred as "check-in") that can also be shared to other social networks such as Facebook or Twitter.

In Foursquare, users are encountered to be very specific with their check-ins indicating their precise location or activity while at a venue, and they receive awards as incentives for checking in. This enables Foursquare to collect important information from the users that can be used to provide personalized recommendations and business deals. Although the real-time location of the users is not shared on the Foursquare app, all the user interaction with a venue, such as writing a tip is time-stamped and publicly available to the community. This allows other users to infer when the user was at a specific place. Moreover, information like checking-in at a place might not be public, but it can be accessed by the followers of the user and allows them to know when the user visited the venue.

Foursquare assumes that the users are aware regarding the location privacy issues and responsible for the visibility of their geo-located posts. Hence, the users can only protect their privacy by defining the visibility of their shared posts.

[2] https://about.twitter.com/company.

[3] https://foursquare.com/about.

14.2.2 Transportation Applications

Another type of apps that require the location of the users are transportation apps. These applications take advantage of the mobile sensors in order to provide real-time navigation to the users. Hence, transportation applications need to acquire the user real-time location frequently to propose the optimal route and to constantly provide directions to a user to reach her destination.

14.2.2.1 Google Maps

Google Maps (http://maps.google.com) is an online service that offers satellite imagery, street maps, panoramic views of streets, real-time traffic conditions and route planning for traveling by foot, car, bicycle or public transportation. Moreover, Google Maps provides an online service for the users to navigate to places that requires sharing the real-time location of the user, captured through the sensors of her mobile device.

Besides navigation purposes, Google Maps also collects the speed and location information of the users anonymously. This is used to calculate traffic conditions in real-time so as to provide better estimations for the travel times. Hence, the users need to accept this fact in order to be able to take advantage of the Google Maps.

Google has mentioned that it permanently deletes the start and end points of every user trip it monitors so that information about where each user came from and went to remains private. Moreover, users can set locations that user activity will not be captured such as their home or work.

14.2.2.2 Moovit

Moovit (http://www.moovitapp.com) is a public transit app and mapping service that features trip planning, real-time arrival and departure times, line schedules, alerts, and advisories that may affect the trip of the user. Moovit uses several transit modes including buses, ferries, metro, trains, trams, and trolleybuses.

Moovit allows users to send reports actively including reasons for delays, overcrowding, satisfaction with their bus driver, and wifi availability which is shared with the community. Moreover, the users can also share data passively. By riding with the 'Live Directions' feature active allows Moovit to collect passively and anonymously the user speed and location data. This data is used in addition to the public transit schedules to improve trip plan results based on current conditions.

Moovit utilizes the information gathered from the application, including information regarding the user location and public transport preferences. They also make anonymous, statistical use of this information in order to estimate the arrival times of various bus lines, analyze their frequency and convey the information to third parties for whom this information is likely to be relevant.

14.2.2.3 Waze

Waze (http://www.waze.com) is a navigation application that provides turn-by-turn information to the users. The Waze app captures real-time information that translates into traffic conditions and road structure from its users. Moreover, it enables them to actively report to the community traffic, accidents, police traps, blocked roads, weather conditions and much more. Waze collects and analyzes this information to provide other Wazers with the most optimal route to their destination.

Similar to the rest of the Transportation Applications, Waze may use anonymous, statistical or aggregated information, including anonymous location information. Any other content that the users submit manually, such as geo-located reports, as well as their current location during a route is publicly available to all users of Waze. However, Waze users are able to share information as anonymous users to preserve their location privacy, that prevents them from collecting the rewards for their contributions. We also note, that, during navigation Waze does not share data within 500 m from the user home to preserve user privacy.

14.2.2.4 Uber

Uber (http://www.uber.com) provides a service for hiring a private driver. It allows users with smartphones to submit a trip request, which is automatically sent to the nearest Uber driver, alerting the driver regarding the location of the customer. Uber drivers typically use their own personal cars and their payment is calculated by the Uber app.

When someone uses Uber her precise location data about the trip is collected from the Uber app used by the Driver. Moreover, if the user permits the Uber app to access location services of the mobile device Uber also collects the precise location of the device when the app is running in the foreground or background.

14.2.3 Applications for Traveling

Traveling applications is another category of apps that share interactive travel-related content including ratings and experiences for specific points of interest. Although users may opt-out from sharing their location with the app, sharing their experiences implicitly validates their presence.

14.2.3.1 TripAdvisor

TripAdvisor (http://www.tripadvisor.com) is one of the early adopters of user-generated content that allows users to contribute content based on places that they

visit to advise the rest of the community based on their experience. TripAdvisor offers advice from millions of travelers for a large variety of places including accommodations, restaurants and attractions. Nevertheless, all these data are implicitly geo-located and shared publicly with the community.

TripAdvisor may collect information about the user location if the user has instructed the device to send such information to the application or if the user has uploaded photos tagged with location information. This location information is used to provide relevant content and contextual advertising to the user.

14.2.3.2 Yelp

Yelp (http://www.yelp.com) uses a similar model with TripAdvisor where users provide reviews and ratings for points of interest that can be exploited by the community when they make decisions. In addition to reviews, Yelp can also be used to find events, lists and to talk with other users. In Yelp every business owner can setup a free account to post photos and messages from their customers. Nevertheless, similar to TripAdvisor, these reviews can reveal the users' spatiotemporal presence.

14.2.4 Fitness Applications

Several fitness applications have recently emerged due to the massive prevalence of smartphones and wearable devices such as smartwatches and smartbands. These devices are equipped with sensors such as GPS, accelerometer, gyrometer, pedometer and heart rate sensor that give the opportunity to the users to capture and share several aspects of their training including distance travelled, steps, heart rate, calories burned, etc. Nevertheless, these applications capture massive amounts of data from the users, especially when the users are active.

14.2.4.1 Strava

Strava (http://www.strava.com) is one of the most prominent fitness apps especially for cyclists. It allows them to record their bike rides and runs, compare their performance over time and share them with the community so that they compare their performance with other users. The users can either share these data with the community to compete with other users or preserve them in their account. However, inevitably, users that share these data allow potentially untrusted users to access their trajectories, including their commutes. In order to preserve user privacy for users that share their data, Strava hides all user traces within a predefined radius from the user's home.

14.2.4.2 Endomondo

Another similar social fitness app is Endomondo (http://www.endomondo.com), which allows users to track their fitness and health statistics with a mobile app. Endomondo is focused on running and walking and encourages users to track their workouts so as to reach their fitness goals. Similar to Strava it allows its users to share their workouts with the community and, thus, share their mobility traces.

14.2.5 Image Sharing Applications

Another type of social location-sharing applications are Image Sharing Applications. These apps allow users to share images with their social interactions or publicly. In addition they also give the opportunity to the user to tag the location where an image is captured.

14.2.5.1 Instagram

Instagram (http://www.instagram.com) is one of the most prominent location sharing applications with more than 500 million users.[4] It allows users to share images with geo-location with the community and interact with the shared images. In Instagram anyone can view the images shared by the users by default. However, each individual user can choose to make them private so that only approved followers can see them.

14.2.5.2 Flickr

Flickr (http://www.flickr.com) provides a similar service for the users to upload their photos and share them with the community and has been widely used by photo researchers and by bloggers. In Flickr, each photo can be geo-located either by the mobile device or manually by the user and the photo's location can then be shown on a map. In Flickr each user can set her geo-privacy to determine the users who are allowed to access the location of the published photos.

14.2.6 Friend/Family Finder Applications

Another type of location sharing applications are the Friend/Family Finder Applications. These applications have been initiated due to the simplicity of acquiring

[4] https://www.instagram.com/about/us/.

and sharing user location in mobile devices. Hence, they track the users in order to notify them when there is a nearby friend and to allow their friends and family to know where they are.

14.2.6.1 GPS Phone Tracker

One of the most popular Friend Finder application is the GPS Phone Tracking Pro App (http://gpsphonetracker.org). It has two main features: (1) it allows users to find their friends and navigate towards them and (2) it allows the user to find her mobile device from the app's website.

GPS Phone Tracker uses the geographical coordinates of the user to report her real-time whereabouts to her friends. Once registered to the system, each friend appears as a unique icon on the map to let the user know where each of her friends is. However, this implies that the users are constantly tracked by the operator but also from all their relatives.

14.2.6.2 Family Locator: GPS Tracker

Another popular app for location sharing is the Family Locator app (https://www.life360.com/family-locator/). It allows users to create their own groups of people such as friends or family, called "Circles". This enables them to view the real-time location of Circle Members and receive real-time alerts when circle members arrive at or leave destinations. Moreover, similar to the GPS Phone tracker, it enables users to locate their mobile devices. Again, the issue is that the users are constantly tracked and their location is shared, providing no privacy among users in a particular circle.

14.2.7 University Initiatives

The following apps are differentiated from the above categorization because they have been initiated by Universities. The main difference of these applications is that they have tried to incorporate innovative techniques to provide users with features that require their real-time location, while preserving user privacy at the same time.

14.2.7.1 PCube

PCube (http://www.everywaretechnologies.com/apps/pcube) is a location-aware social networking app that aims to alert users when their friends are in close proximity. Hence, it allows users to observe which of their friends happen to be in the area.

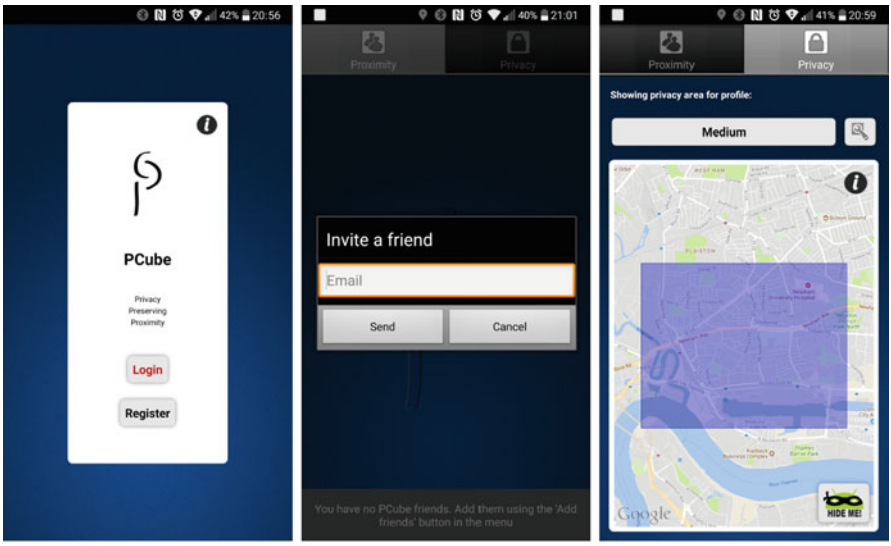

Fig. 14.1 The PCube app. (**a**) Login. (**b**) Add friends. (**c**) Set proximity

As we mentioned above, such Friend/Family Finder apps need to constantly share the user location to provide this functionality. However, PCube aims to preserve user privacy for the location-based data shared by the users using a novel approach, further described in [18]. Hence it exploits encryption techniques to encrypt the data before leaving the device to prevent user exposure even from their servers. Moreover, it takes advantage of the privacy preferences, set by the users, to alert friends in proximity, but the alert received does not reveal the actual user location with higher precision than the one she has decided.

Figure 14.1a illustrates the login screen of PCube that enables users to login and register. After login, the user is presented with a screen which is divided into two tabs. The first tab (Fig. 14.1b) allows the user to add friends and create social communities while the second one (Fig. 14.1c) provides the proximity functionality to set the radius within which her friends will be alerted that they are in close proximity.

14.2.7.2 Locaccino

Another Friend/Family Finder app initiated by the academic community, is Locaccino (http://locaccino.org), that allows users to share their location with their friends.

Similar to PCube, Locaccino focuses on preserving user privacy but it uses a different approach. Locaccino exploits a rule-based approach that allows users to create rules rules regarding their visibility from the privacy settings. The rules can be as simple or complex as the user wants them to be, such as setting a different

Fig. 14.2 The Locaccino app. (**a**) Map. (**b**) Privacy settings. (**c**) History

proximity for each particular user or disabling the visibility of the user for specific hours or places.

Figure 14.2a shows the map presented to the user that presents the geo-location shared by her acquaintances. Moreover, the user can modify the rules that will be used to enable others to access her location, as shown in Fig. 14.2b, and the view the history of the users that could access her location, as illustrated in Fig. 14.2c.

14.2.7.3 CrowdAlert

CrowdAlert (http://crowdalert.aueb.gr) is a location-based app, designed for Android users that enables them to report real-world events of interest and receive real-time alerts for nearby events, by exploiting real-time data from human users, road sensors, bus sensors and web sources. Similar to Social and Transportation apps it enables users to provide sparse geo-located data such as geo-located reports for real-world events (e.g., Traffic, Floods, etc). Such data can reveal the location of the user to the community. However, CrowdAlert is built with privacy in mind and focuses on effectively preserving user privacy for her shared data as we describe in detail in Sect. 14.5.

14.3 Privacy Characteristics of Location Sharing Applications

In this section we present the types of data shared by the users in location-based applications. Then we explain the information that can be inferred by the data and finally we describe thoroughly the potential privacy risks that arise from exposing such data per service.

14.3.1 Type of Data Shared

Location-based applications share different types of content but in all applications the content is coupled with the user location that can lead to privacy implications. The data shared in each application is illustrated in Tables 14.1 and 14.2. In the following we enumerate the most important categorization among the different types of data sharing in terms of privacy.

14.3.1.1 Data Density

One important categorization depends on the density of the sharing location-based data. Applications such as Transportation applications, Fitness applications and Friend/Family Locator applications typically share dense data that include the user mobility, such as user trajectories. On the other hand, Social applications, Travel applications and Image Sharing applications typically share sparse data since every post is coupled with a specific geographical location.

Obviously, providing dense data makes the privacy preservation task even more difficult as an adversary acquires richer information from the system. However, it has been proved that even a small number of spatiotemporal points is enough for user identification [7]. Hence, even sparse data might be enough for an adversary to expose user privacy.

14.3.1.2 Visibility

We can also discriminate the applications depending on whether they allow sharing the location-based data publicly. For instance, Social, Image Sharing and Fitness Applications give the opportunity to the users to share their data publicly if they want to. On the contrary, the majority of the Transportation apps allow users only to share data privately with the service operator or anonymously. Based on this characteristic we argue that public data are prone to exposing user privacy since all the data are linked with a specific user identifier. Furthermore, even anonymized data can be de-anonymized through complex data analysis [10].

14.3.1.3 Data Source

Finally, another categorization of the data is whether they are generated from the users explicitly or they are generated by the mobile sensors. Hence, Social, Image Sharing and Travel applications require from the users to generate the content which is coupled with the geographical location. On the other hand, Transportation applications and Friend/Family Locator applications produce dense data which is extracted from the mobile sensors without requiring user interaction.

Table 14.1 Location-based applications, data sharing and privacy mechanisms

App	Type	Google play downloads	Data shared	Privacy mechanism
Facebook	Social	1,000,000,000–5,000,000,000	Share posts with the community Posts may include geo-location	Visibility scopes
Twitter	Social	500,000,000–1,000,000,000	Share tweets with users Tweets may include geo-location	Visibility Scopes
Foursquare	Social	10,000,000–50,000,000	Share user attendance with friends (check-in) Share public content like public user profile information, tips, likes, saves, public photos, badges/stickers, mayorships, and lists of friends	Visibility scopes
Yelp	Travel	10,000,000–50,000,000	Users can share reviews regarding places that include implicit location	Disabling location sharing
TripAdvisor	Travel	100,000,000–500,000,000	Users can share reviews and ratings regarding places that include implicit location Users can share their location with TridAdvisor to receive relevant content and contextual advertising	Disabling location sharing
Moovit	Transportation	10,000,000–50,000,000	Users can passively and anonymously transmit their speed and location Users can actively provide reports including reasons for delays, overcrowding, etc	Anonymous data sharing
Waze	Transportation	100,000,000–500,000,000	Users can share geo-located traffic reports that may be viewed by all users in the community Waze may use anonymous, statistical or aggregated information	Privacy zones Anonymous data sharing
Google traffic maps	Transportation	1,000,000,000–5,000,000,000	User accurate location is shared during navigation Users share live traffic conditions anonymously (speed and location information)	Privacy zones Anonymous data sharing
Uber	Transportation	100,000,000–500,000,000	Share user location with drivers, other riders, general public, third parties	Disabling location sharing
Strava	Fitness	5,000,000–10,000,000	Users shares bike rides, runs Users can also share their achievements	Privacy zones Visibility scopes

(continued)

Table 14.1 (continued)

App	Type	Google play downloads	Data shared	Privacy mechanism
Endomondo	Fitness	10,000,000–50,000,000	Users can shares their workouts Users can also share their achievements	Visibility scopes
Instagram	Image sharing	1,000,000,000–5,000,000,000	Share images with geo-location	Visibility scopes
Flickr	Image sharing	10,000,000–50,000,000	Share images with geo-location	Visibility scopes
GPS phone tracker	Friend finder	10,000,000–50,000,000	User shares her real-time location with her friends	Disabling location sharing
Family locator—GPS tracker	Friend finder	10,000,000–50,000,000	User shares her real-time location with her friends	Disabling location sharing

Table 14.2 University initiative location based app applications, data sharing and privacy mechanisms

Apps	Type	Data shared	Privacy mechanism
PCube	Friend finder	Alerts the users about friends who are in proximity	Rule-based scheme encryption
Locaccino	Friend finder	Share user location with friends	Rule-based scheme
CrowdAlert	Event alerting	Share public geolocated reports with the community	Data suppression

14.3.2 Inferring Information from the Data

The problem of exposing all this data is that they reveal the user's spatiotemporal instance when the data are published. Assume a user that passively observes and records the location information published by an individual user. Although the information shared in a particular post may not reveal a lot of information when considered individually, it can expose a large fraction of the user mobility as the user shares more and more data with the community. The problem can become worse when the user identifier can be linked across different applications and, thus, all geo-located information shared across these networks can be linked to the specific user.

The issue with sharing location-based data is that the majority of the data typically reside at locations or along routes that are mostly visited by the individuals. Hence, as the amount of content contributed by the user increases, the shared data will gradually expose the user's most frequent trajectories, important locations and

even the user's physical identity, which could place a user in physical danger [15] and lead to crimes such as stalking the user or robbing the user when she is absent from home.

We also note that due to the nature of the data, an adversary might be able to generate the user mobility even using sparse data. For instance assume that a user produces a couple of geo-located posts in a social app as she commutes from her home to work. Although these data can be very distant, the intermediate points can easily be inferred using navigational tools. Furthermore, such information can also be linked with posts shared previously at the same time window where the user commutes to her work, in case that they overlap spatially.

14.3.3 Privacy Threats

The data shared in such applications enable an adversary to pose the following three types of privacy threats [9]: (1) Tracking Threat, (2) Identification Threat and (3) Profiling Threat. These threats originate from sharing location-based data with potentially untrusted entities and need to be taken into consideration when developing an application that shares the user location to be able to preserve user privacy.

- **Tracking Threat:** In a location-based application an untrusted party might extract continuous location updates that enable him to track the user in real-time. Hence, an adversary should not be able to determine the user mobility and predict her future location with high accuracy when leveraging the location-based data shared by the user.
- **Identification Threat:** Even if the untrusted party is able to sporadically accesses a user's location, the untrusted party should not be able to identify the user's most frequently visited places, such as the user's home and work location. This is because an adversary can exploit such information to reveal the identity of the user even from anonymous mobility traces.
- **Profiling Threat:** The user mobility traces, shared by the mobile application, might not reveal only places that can help to identity the user but also places that can be used by an adversary to profile the user. For instance, an adversary will be able to profile the user when he acquires location data showing that the user has visited religious places or attended political meetings.

Note, that, the user might explicitly share information that reveal sensitive information (e.g., posting a geo-located tweet that he just arrived at home). Although a privacy-concerned app will not prevent the user from sharing this information, it should consider these privacy threats when this information can be inferred from the shared data.

14.4 Existing Approaches for Privacy Preservation

14.4.1 Privacy Mechanisms in Popular Location-Based Applications

Subsequently, we present the mechanisms which are used in each popular location-based service to prevent unauthorized parties from learning the user's current or past locations. The main schemes can be grouped into the following categories:

1. **Visibility Scopes** allow the user to select the groups of users which are allowed to access the user's posts and respective location.
2. **Privacy Zones** essentially hide the user location in specific areas to anonymize the user.
3. **Rule Based approaches** where the users develop rules per user group in order to determine the visibility of their data depending on multiple factors.
4. **Anonymous Data Sharing** allow the users to share location based data but the service consider them as anonymous data, without coupling them with the user.

14.4.1.1 Visibility Scopes

The most prominent approach used in popular location-based services is the use of visibility scopes. This approach enables the users to develop groups of users such as friends and family manually. Then they can select for every piece of information that they share if the information will be publicly available or determine the groups of users that will be able to access this information. The idea behind visibility scopes is that the users are capable of managing the visibility of their posts and evaluate their privacy exposure.

Visibility scopes have been widely used as a privacy preservation paradigm in multiple social location-based services such as Facebook, Twitter and Foursquare. This is because users in these networks tend to share content with multiple users and thus they do not focus on their privacy. Similarly, Image Sharing applications like Flickr and Instagram depend on the users to decide who can view their photos as well as the location of their photos.

However, this approach can lead to serious privacy issues especially when the shared data embed the users' geographical location. The problem with visibility scopes is that they require constant effort from the user to preserve her privacy. Hence, users tend to share information publicly or with users that they have accepted as their friends, that they might have never met. Thus, as we mentioned above sharing this information can be exploited by the community to track the user, extract sensitive user information or profile the user.

14.4.1.2 Privacy Zones

Another well-known approach that also requires user interaction are Privacy Zones. This approach allows the user to select one or more sensitive locations that will be

considered by the app in order to hide information produced near these locations. This is typically achieved using a radius around sensitive locations within which no information will be shared with the community.

Privacy Zones have been used particularly for apps that share frequent user locations such as the Strava fitness app and Waze. Hence, in Strava when a user shares a bike route the data near the sensitive locations will be hidden from the rest of the users. Similarly, Waze hides all the data within 500 m from the user's home.

Although Privacy Zones are more sophisticated than Visibility Scopes as they require less interaction with the user, they have one important drawback. Since the radius is typically fixed in these apps, an adversary can easily identify the sensitive locations with high accuracy with triangulation. However, even if the radius can be modified, the sensitive locations will start to be exposed as the users contribute more data.

14.4.1.3 Rule-Based Scheme

A different technique to preserve user privacy is the Rule-Based Scheme. This approach also requires user interaction in order to determine several rules regarding her visibility. However, these rules can be very complex, such as hiding the user location for specific places, hours or people (e.g., hiding the user location from her employer during work hours).

Rule-Based approaches have been used from Friend Finder applications such as PCuble and Locaccino, that produce massive location-based data from the users as they constantly track them. Hence, it allow them to filter the portion of the data that each individual will be able to observe.

Rule-Based approaches constitute an improvement compared to Visibility Scopes and Privacy Zones as they use their concepts but extend them to make more complex decisions regarding the data that can be shared among users.

14.4.1.4 Anonymous Data Sharing

Another technique employed by the commercial location-based applications to protect privacy for privacy-concerned users is Anonymous Data Sharing. This approach allows users to share data with the service provider or the community without revealing their identifiers.

Anonymous Data Sharing has been extensively used for traffic and transportation related applications such as Moovit, Google Maps, HERE Maps and Waze. These applications collect user data from mobile sensors, typically using a background service, to compute real-time road conditions. Additionally, user-identified services can also be provided in an anonymous way. For instance, Waze facilitates anonymous contributions for privacy-concerned users. However, such alternatives are not always attractive to end users, as they do not receive any reward for their contributions, and from the system's point of view, anonymous data has shown to increase spamming.

14.4.1.5 Disabling Location Sharing

Finally, another option for the users to preserve their privacy is to completely turn off the "Location Services" from their devices. This action prevents the mobile operating system to give access to the location-based readings of the mobile sensors to the applications. Similarly a user can typically reject permission to access the location data for a particular app. This approach can be used by the users for applications like Yelp, TripAdvisor, Uber, GPS Phone Tracker and Family Locator - GPS Tracker, as they do not provide any other way to prevent tracking the mobility traces of the user. However, this option is out of the scope of this chapter since it removes all the functionality that depends on the user location, making some of these applications ineffectual.

14.4.2 Other Approaches

Finally, several approaches have been proposed in the literature to deal with privacy threats and preserve user privacy for location-based application. We summarize them here for completeness but they are discussed in detail in Chap. 5.

The first set of privacy preservation approaches focuses on modifying location-based information using a variety of mechanisms: (a) **Path Confusion** [8, 14] that aim at confusing the paths of users that reside in the same region by connecting their traces, (b) **Mix Zones** [11], which are generated when there are enough users located in the same place at the same time, making it hard for an adversary to distinguish an individual from others that reside in the same zone at the same time, (c) **Fake Data Injection** [12, 20], where fake data are injected along with real data to confuse an attacker, (d) **Data perturbation** [2, 25] where the goal is to modify the original data set with some noise drawn from a selected statistical distribution to preserve user privacy, (e) **Data generalization** [9, 17] that generalizes the user location information by reducing the spatial accuracy of the user, and (f) **K-anonymity approaches** [6, 16, 19] that release data that hold the k-anonymity property.

The second set of approaches to protect user privacy focuses on maintaining the original location-based data in a way that an adversary will not be able to extract sensitive user information. Here we have two main categories: (a) **Encryption techniques** [13, 23] that encrypt the user data before sharing them so that they disable untrusted parties from accessing the data, and (b) **Data Suppression techniques** [5, 24] where the user location is suppressed when the set of data shared by the user can expose her privacy. That way they prevent adversaries from exposing sensitive user information although they are capable of publishing the original information shared by the users.

14.5 CrowdAlert: Experiences from Deploying a Location-Based App in a Smartcity

In this section, we discuss our experience from deploying CrowdAlert, a location-based application, in a Smart City environment [4]. We explain our design objectives for our privacy preservation approach and we discuss the trade-offs between privacy and utility when publicly sharing data. Moreover, we give an overview of how our privacy approach works and we point out the effort needed to develop an interface that can satisfy both novice and expert users. Finally, we discuss our experience and the knowledge that we obtained from deploying CrowdAlert.

CrowdAlert is a location-based mobile application that enables users to report and receive information regarding unusual events in SmartCities that include Accidents, Constructions, Traffic, Natural Disasters, etc. CrowdAlert provides great benefits to both citizens and authorities in a SmartCity. In particular, it allows citizens to be alerted about local unusual events in real-time, and gives the opportunity to city authorities to identify, supervise and react if necessary, to these events in a cost-effective manner.

14.5.1 Data Sharing in CrowdAlert

Users connect with CrowdAlert using the mobile application, which is freely available at Google Play.[5] All registered users receive real-time information for a wide variety of events such as Accidents, Constructions, Hazards and more. This information arrives both in the form of notifications as well as on a map-based interface as can be seen in Fig. 14.3a. Moreover, registered users are able to report such information through CrowdAlert, as can be observed in Fig. 14.3b and receive questions for local information from the authorities when needed.

CrowdAlert app is designed to consider user privacy preservation, since users share information, coupled with their personal geographical location, that can lead to privacy exposure. Hence, we have incorporated our privacy preserving mechanisms to allow users share data with the community without compromising their privacy, as we explain in the following.

14.5.2 Design Objectives for Data Sharing

Before developing our privacy preservation approach we set the following design objectives for our solution:

[5]CrowdAlert—http://crowdalert.aueb.gr/.

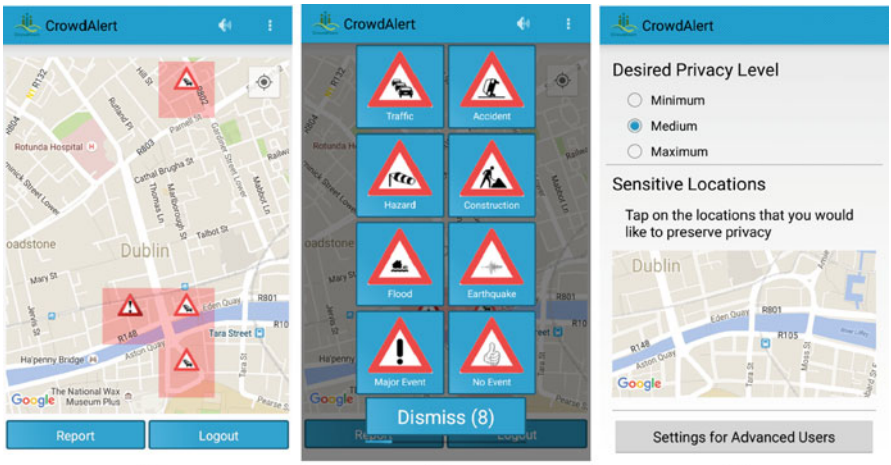

Fig. 14.3 The CrowdAlert app. (**a**) Login. (**b**) Report. (**c**) Settings

1. **Users should be able to post data coupled with their identifier, location and timestamp.** In CrowdAlert we allow the users to share data publicly with their identifiers, location and timestamp as happens in other real-world applications like Social and Transportation applications. However, users are not capable of estimating the effect of the sharing such data on their privacy. The goal of our approach is to alert the user before sharing location-based data that may expose her privacy to decide whether to share it.
2. **Users should be able to tune their privacy levels.** Every user has a different preference in terms of privacy levels and desired utility. However, existing approaches are basically limited to only two options. Users can either allow the location-based application to have access their data or accept generic, potentially low quality levels of service such as anonymous contributions that does not allow the users to receive any reward for their contributions. Our goal is to allow the users to tune the privacy levels they aim to achieve when contributing data.
3. **Users should keep their mobility traces locally.** The majority of the privacy protection approaches that exist in the literature [3], use a centralized approach to analyze user mobility traces. However, such centralized approaches have important shortcomings: (1) they are prone to different types of attacks like eavesdropping and the possibility that the server might become malicious, (2) they introduce additional communication costs that may degrade the user experience that needs to wait for a response before sharing her data. Our goal is to be able to achieve privacy preservation by evaluating the privacy exposure of the user locally on her mobile device and share the data after deciding that they pose no privacy risks.

14.5.3 Trade Offs Among Privacy and Sharing

Due to our first design objective, it is clear that existing solutions to preserve user privacy, that modify the original user data, could not be used. Since we focus on data which is publicly available we used an approach that proactively determines whether sharing information can expose user privacy and suppress such data.

The challenge is that users share content publicly coupled with their spatiotemporal instance as happens in Social Apps or Transportation Apps. Hence, users share more data in different locations, they expose larger parts of their mobility. In CrowdAlert we deal with the privacy preservation problem using a novel data suppression approach that aims to balance the trade off among user privacy and amount of shared data, based on the user preferences.

14.5.4 Privacy Problem

The problem that we deal with is how to prevent honest-but-curious adversaries that monitor the shared geo-located, user identified data from extracting sensitive information about the users. We define that the set of shared data effectively preserves user privacy only if:

- An attacker cannot approximate the user trajectories, and
- The user sensitive locations such as home and work cannot be determined by an adversary from the shared data based on their amount and frequency.

Essentially, our goal is to answer effectively the following questions:

1. How to evaluate the privacy exposure of the user trajectories from the set of the publicly shared reports on resource-constrained mobile devices in real-time?
2. How to preserve user privacy in terms of the user's sensitive locations when sharing geo-located data publicly?
3. Should the user share a newly produced report without the risk of compromising her privacy?

14.5.5 Privacy Preservation Approach

To preserve user privacy for CrowdAlert we developed our privacy mechanism called PROMPT [5]. The core idea of PROMPT is to employ the novel geometric approximation approach of ϵ-Coresets. We selected the ϵ-Coreset approach since ϵ-Coresets can effectively reduce the computation overhead for several complex geometric and graph problems [1]. Thus, they can be used to effectively process user mobility in resource constrained devices. The idea of ϵ-Coresets is to select a small subset that approximately represents the original data that is able to process a given

query with up to $(1+\epsilon)$-multiplicative error. This reduction allows us to perform queries on the coreset and greatly reducing the computation time.

PROMPT is executed whenever a user desires to publish content with the community. It exploits the user mobility traces, which are compressed and preserved locally in the mobile device, in order to compute whether user privacy can be exposed. Hence, we extract all user trajectories that reside near the current user location and then we develop a coreset for each of these trajectories. Each coreset is developed by the current user location and the set of locations that the user has previously shared which are spatiotemporally close to the trajectory. This allow us to compute the fraction of the trajectories that can be approximated by the shared user data and prevent exceeding a predefined threshold, which can be modified by the user, in order to tune her privacy levels. Hence, whenever this threshold will be exceeded we alert the user, before sharing the data, to prevent her from exposing her privacy.

Moreover, in PROMPT we aim to preserve privacy near the user sensitive location; sensitive locations can be determined automatically as the most frequent locations based on the user mobility which is available on the mobile device and manually from the user. However, as mentioned above Privacy Zones are prone to expose sensitive locations through triangulation. Hence, in PROMPT we allow users to share data near sensitive locations but we aim to prevent the user from sharing a large percentage of location-based data near her sensitive locations, compared to the rest of the locations, to confuse an adversary regarding their importance. This can be achieved using the entropy metric [22]. Thus, in PROMPT, we only allow users to share data near sensitive locations when the entropy of the shared data increases. Since entropy increases when the frequency of the locations becomes more similar, we discourage sharing a report near a sensitive location when there already exists a large number of location-based data contributed from that location, compared to other locations.

14.5.6 Privacy Settings Interface

PROMPT has been integrated in a beta version of CrowdAlert to preserve user privacy. The user interface to adjust the privacy levels of PROMPT is illustrated in Fig. 14.3c. Since PROMPT depends on several variables, in order to allow the users to tune their privacy levels, we allow them to select their desired privacy level with three privacy profiles: Minimum, Medium and Maximum. This is due to several recent surveys which state that the users often find it difficult to adjust their privacy preferences [21]. Hence, profiles are able to simplify the privacy choices for the users [26].

Nevertheless, we give the opportunity to advanced users to access the parameters of our approach and fine tune their privacy. An advanced users can adjust: (1) the maximum percentage approximation can be achieve for her trajectories before preventing the user to share data, (2) the spatial radius to consider from the user

sensitive locations to preserve privacy, and (3) the spatiotemporal thresholds that are used to correlate similar user presences.

Moreover, as can be observed from Fig. 14.3c the users can also define their sensitive locations using a map-based interface. These locations are considered from PROMPT in addition to the most frequent locations extracted from the local mobility traces.

14.5.7 Experiences from Our Deployment and Lessons Learnt

The development of CrowdAlert enabled us to interact with best-testers and with the end-users that gave us important feedback which led to the re-design of several aspects of our approach in terms of privacy preservation. There were three significant lessons that we learned from deploying our CrowdAlert app in the real-world that we discuss in the following.

14.5.7.1 Privacy Preservation

Our original version of CrowdAlert was developed without any privacy mechanism as we expected users would be sharing location-based data with multiple services. However, when we shared the app with our beta-testers this became an issue because the app required the users to share their location as a permission upon installation. During our discussions we understood that, although, the app acquired the user location when the user wanted to share a public report, it was not clear to the users whether their privacy could be exposed. Moreover, we felt that there were some users that were extremely concerned regarding their privacy.

In order to deal with this problem we developed our PROMPT privacy preservation approach that was briefly presented above. PROMPT enables users to quantify their privacy exposure in terms of location before sharing their data. Employing such a mechanism made the users were more willing to share their data with the community.

Furthermore, as mentioned above, our PROMPT approach enables the users to tune their privacy levels. This is essential to the privacy preservation mechanism since we realized that different users had different needs. That way we can attract both privacy concerned users as well as users that aim to share information with the community without prioritizing their privacy.

14.5.7.2 Data Sharing

Another important lesson was that the privacy preservation should not degrade the utility of the application. We explored several existing solutions before developing PROMPT, but none of them was able to preserve user privacy effectively without

degrading system utility. The majority of the existing approaches modify user information and, thus, were rejected since our app focuses on real-world events that require accurate location. On the other hand since we focus on public data reports, similar to Waze, an encryption-based approach would not provide any benefit to the user privacy.

To deal with this problem we developed PROMPT, which is based on data suppression to preserve user privacy. One significant benefit of PROMPT, as we show in [5] is that it can preserve high levels of utility in terms of data sharing even with strict privacy settings. This enables even privacy concerned users to be able to share a lot of data without exposing their sensitive information.

14.5.7.3 User Interface

Finally, the most important lesson was that the user interface can play a fundamental role to the users' understanding of their privacy exposure. Our initial interface allowed the users to modify several parameters of our PROMPT approach. Nevertheless, this freedom was also a bottleneck especially for novice users did not have the technical knowledge to understand and set these parameters. Hence, several beta testers complained regarding the privacy settings as they were not able to tune them properly.

In order to deal with the above problem, we developed a new user interface where the privacy setting was selected among three privacy profiles, as explained above. This feature enabled the users to select their desired privacy levels even if they did not tune every aspect of the algorithm. Nevertheless, we kept the parameter tuning as a separate functionality that can be accessed by advanced users.

14.6 Conclusion

In this chapter we provided a survey of the most popular location-based applications and we explained the benefits that they provide to the users. Then, we discussed the respective privacy threats that may arise in such settings and presented the existing privacy mechanisms. Finally, we presented our own experiences upon developing a real-world location-based application, the privacy mechanism that we developed and the lessons that we learnt based on the issues that we encountered.

One important message that can be extracted from this chapter is that location-based applications can expose a great deal of sensitive user information and that the most popular applications rely on the user to handle such issues. Although, several approaches have been proposed from the research community, the commercial applications use simple solutions and expect that the users are capable of quantifying their exposure and preserve their privacy. On the other hand, we argue that popular real-world applications should benefit from the solutions that have already been proposed in the research literature so as to allow users to continue contributing information without being concerned about their privacy.

References

1. P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Geometric approximation via coresets. *Combinatorial and computational geometry*, 52:1–30, 2005.
2. R. Agrawal and R. Srikant. Privacy-preserving data mining. In *SIGMOD*, Dallas, Texas, United States, May 2000.
3. C. Bettini and D. Riboni. Privacy protection in pervasive systems: State of the art and technical challenges. *Pervasive and Mobile Computing*, 17:159–174, 2015.
4. I. Boutsis and V. Kalogeraki. Crowdalert: a mobile app for event reporting and user alerting in real-time. In *UbiComp*, Heidelberg, Germany, 2016.
5. I. Boutsis and V. Kalogeraki. Location privacy for crowdsourcing applications. In *UbiComp*, Heidelberg, Germany, 2016.
6. C.-Y. Chow, M. F. Mokbel, and W. G. Aref. Casper*: Query processing for location services without compromising privacy. *ACM Transactions on Database Systems (TODS)*, 34(4):24, 2009.
7. Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports*, 3, 2013.
8. K. Dong, T. Gu, X. Tao, and J. Lu. Complete bipartite anonymity: Confusing anonymous mobility traces for location privacy. In *ICPADS*, pages 205–212, Singapore, December 2014.
9. K. Fawaz and K. G. Shin. Location privacy protection for smartphone users. In *CCS*, pages 239–250, Scottsdale, Arizona,, november 2014.
10. S. Gambs, M.-O. Killijian, and M. Núñez del Prado Cortez. De-anonymization attack on geolocated data. *Journal of Computer and System Sciences*, 80(8):1597–1614, 2014.
11. S. Gao, J. Ma, W. Shi, G. Zhan, and C. Sun. Trpf: A trajectory privacy-preserving framework for participatory sensing. *Information Forensics and Security, IEEE Transactions on*, 8(6):874–887, 2013.
12. A. Gkoulalas-Divanis and V. S. Verykios. A privacy-aware trajectory tracking query engine. *SIGKDD Explorations Newsletter*, 10(1):40–49, May 2008.
13. T. Higuchi, P. Martin, S. Chakraborty, and M. Srivastava. AnonyCast: privacy-preserving location distribution for anonymous crowd tracking systems. In *UbiComp*, pages 1119–1130, Osaka, Japan, Sep 2015.
14. B. Hoh and M. Gruteser. Protecting location privacy through path confusion. In *SECURECOMM*, Athens, Greece, September 2005.
15. C.-C. Hung, W.-C. Peng, and W.-C. Lee. Clustering and aggregating clues of trajectories for mining trajectory patterns and routes. *The VLDB Journal*, pages 1–24, 2011.
16. A. Kapadia, N. Triandopoulos, C. Cornelius, D. Peebles, and D. Kotz. Anonysense: Opportunistic and privacy-preserving context collection. In *Pervasive Computing*, pages 280–297. Sydney, Australia, May 2008.
17. S. Mascetti, L. Bertolaja, and C. Bettini. Safebox: adaptable spatio-temporal generalization for location privacy protection. *Transactions on Data Privacy*, 7(2):131–163, 2014.
18. S. Mascetti, D. Freni, C. Bettini, X. S. Wang, and S. Jajodia. Privacy in geo-social networks: proximity notification with untrusted service providers and curious buddies. *The VLDB Journal*, 20(4):541–566, 2011.
19. B. Niu, Q. Li, X. Zhu, G. Cao, and H. Li. Achieving k-anonymity in privacy-aware location-based services. In *INFOCOM*, pages 754–762, Toronto, CA, April 2014.
20. N. Pelekis, A. Gkoulalas-Divanis, M. Vodas, D. Kopanaki, and Y. Theodoridis. Privacy-aware querying over sensitive trajectory data. In *CIKM*, Glasgow, Scotland, October 2011.
21. N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao. Understanding and capturing people's privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing*, 13(6):401–412, 2009.
22. C. E. Shannon. A mathematical theory of communication. *SIGMOBILE Mob. Comput. Commun. Rev.*, 5(1):3–55, 2001.

23. J. Shao, R. Lu, and X. Lin. Fine: A fine-grained privacy-preserving location-based service framework for mobile devices. In *INFOCOM*, pages 244–252, Toronto, CA, April 2014.
24. M. Terrovitis and N. Mamoulis. Privacy preservation in the publication of trajectories. In *MDM*, Beijing, China, April 2008.
25. I. J. Vergara-Laurens, D. Mendez, and M. A. Labrador. Privacy, quality of information, and energy consumption in participatory sensing systems. In *PerCom*, pages 199–207, Budapest, Hungary, March 2014.
26. S. Wilson, J. Cranshaw, N. Sadeh, A. Acquisti, L. F. Cranor, J. Springfield, S. Y. Jeong, and A. Balasubramanian. Privacy manipulation and acclimation in a location sharing application. In *UbiComp*, pages 549–558, Zurich, CH, September 2013.