Privacy through Uncertainty in Location-Based Services

Peer-reviewed author version

Merrill, Shawn; Basalp, Nilgun; Biskup, Joachim; Buchmann, Erik; Clifton, Chris; KUIJPERS, Bart; OTHMAN, Walied & Savas, Erkey (2013) Privacy through Uncertainty in Location-Based Services. In: Proceedings of PriSMO: Privacy and Security for Moving Objects, p. 1-6.

Handle: <http://hdl.handle.net/1942/14913>

Privacy through Uncertainty in Location-Based Services

Shawn Merrill*, Nilgün Basalp[†], Joachim Biskup[‡], Erik Buchmann[§], Chris Clifton[¶],
Bart Kuijpers^{||}, Walied Othman^{**}, and Erkey Savas^{††}

*Department of Computer Sciences

Purdue University, West Lafayette, Indiana 47907-2107

Email: smerrill@cs.purdue.edu

[†]Department of Law

Istanbul Bilgi University, Yahya Köprüsü Sokak No: 1, 34440 Dolapdere, Beyoğlu, TR

Email: nilgun.basalp@bilgi.edu.tr

[‡]Department of Computer Sciences

Technische Universität Dortmund, Otto-Hahn-Str. 20, 44227 Dortmund, DE

Email: joachim.biskup@cs.tu-dortmund.de

[§]Department of Computer Science

Karlsruhe Institute of Technology, Karlsruhe, DE

Email: buchmann@kit.edu

[¶]Department of Computer Sciences and CERIAS

Purdue University, West Lafayette, Indiana 47907-2107

Email: clifton@cs.purdue.edu

^{||}Databases and Theoretical Computer Sciences

Hasselt University, 3590 Diepenbeek, BE

Email: bart.kuijpers@uhasselt.be

^{**}Department of Geography

University Zurich - Irchel, Winterthurerstr. 190, CH-8057 Zurich, CH

Email: walied.othman@geo.uzh.ch

^{††}Department of Engineering and Natural Science

Sabancı University, Istanbul, TR

Email: erkays@sabanciuniv.edu

Abstract—Location-Based Services (LBS) are becoming more prevalent. While there are many benefits, there are also real privacy risks. People are unwilling to give up the benefits - but can we reduce privacy risks without giving up on LBS entirely? This paper explores the possibility of introducing uncertainty into location information when using an LBS, so as to reduce privacy risk while maintaining good quality of service. This paper also explores the current uses of uncertainty information in a selection of mobile applications.

The prevalence of Location-Based Services (LBS) in modern technology has created new risks to user privacy. Since users are unwilling to surrender many of the benefits of LBS, new methods are needed to reduce such risks. This paper explores the possibility of introducing artificial uncertainty into location information when using an LBS without rendering it useless. Uncertainty occurs naturally, so LBS are likely to work in spite of uncertainty. For example, Figure 1 shows location determined by an Apple iPad at Schloss Dagstuhl. Initially, the location was reported with a high degree of uncertainty. Later, the uncertainty was reduced – but the location was not exactly as reported. Our question is: Can we protect privacy by providing uncertain location data, while still retaining good service?

To motivate this problem, let us first recall some instances where these privacy concerns arose. One such case, which received considerable popular attention, is Apple storing and collecting location data from its users' iPhones, unbeknownst to the user. The issue was uncovered on April 20th, 2011. Researchers discovered a file, consolidated.db, that contained longitudes and latitudes combined with a timestamp. This file contained locations that dated as far back as the release of iOS 4, which made it contain a year's worth of location data, stored on the iPhone, synced (backed up) with iTunes and transmitted to Apple, all without the user's knowledge [1].

The "Please Rob Me" website is another prime example of the risks of location privacy (<http://pleaserobme.com/>). The website seeks to raise awareness of the risks posed to location privacy, specifically it provides an interactive map of the places a user checks-in to point out when the user is not at home. The check-in history allows the system to determine the most likely location for the user's home. While our discussion will not include a solution to this issue, the website highlights the prevalence and practical risk posed by lack of location privacy.

Another particularly poignant example is the application "Girls Around Me", which combines social media info and location to find nearby women (who hadn't necessarily said

Fig. 1. Location uncertainty at Schloss Dagstuhl (actual location in hallway outside room N009)

they want to be found) [2]. Additionally, with one click the user can access the Facebook profiles of any of the pictured girls.

There are numerous other programs that collect location data from mobile devices. For the purposes of our discussion, we will only concern ourselves with those applications which collect data and use the information to provide an agreed-upon service. It should be noted that there are also applications which use location information, but not to provide any end-user service, however these apps will not be considered in our discussion. These two ideas are not mutually exclusive, for instance a mapping application that records a timestamp may need longitude and latitude but not the time when you were at that location. It should also be noted that there exist methods for preventing the collection of location data entirely, however these methods completely nullify the benefits of applications which provide location services. Our primary focus, then, concerns discussing methods that hinder the undesired tracking capacities of applications, while still ensuring the same quality from desired services.

The outline of the paper is as follows: In Section I we establish what is meant by privacy for the purposes of this discussion. In Section I-A we discuss the legal status of collecting and using location data by applications. Section II discusses the privacy considerations we will be concerned with here, and potential methods for introducing uncertainty into location data. For the purposes of this discussion, by “uncertainty” we will mean any method which seeks to intentionally decrease the confidence of a third party to determine a user’s specific location. Section III discusses the current impact of uncertain location data on quality of service for applications where location is inherently relevant to the requested service.

I. PRIVACY ANALYSIS

One key challenge that must be addressed is how to analyze privacy. While there has been some research in this area (e.g., [3], [4]), this is still a challenge with room for research, especially with regard to real world applications.

We are considering the following agents: an application *provider* offering one or more dedicated services to a certain class of *clients*, which might be formed by subscription or even on an ad-hoc basis. Each of the clients might repeatedly request one of the services. Each request comes along with data about the requester’s location in order to enable the provider to return a location-dependent reaction to the requester. Without privacy-preserving measurements, the location would be determined in the best possible way, uncertainty is solely a result of technical limitations in determining the precise position. In addition, each request is associated with a time stamp, which reflects real time by default. Introducing additional mechanisms for privacy-preservation enables a client to purposely generate further uncertainty about his or her actual location. Though the time data might be blurred as well, for the sake of simplicity we leave this option open to future considerations. Accordingly, over time the provider receives a sequence of requests, each consisting of at least the following components: location, time, kind of request.

Beyond supposedly providing the service requested honestly, the provider may behave in a *curious* way, *analyzing* the collection of request data received so far for the sake of any secondary use. We discuss four types of secondary analysis that are particularly relevant to location data; while not a complete list, these do give an idea of the kinds of things a service provider may do with location data and the potential impact to users. Moreover, we emphasize the following need of a client: in order to decide about the employment of uncertainty generating mechanisms, a requesting client has to evaluate the potentials for a successful analysis by the provider according

to some metrics.

Location-based reidentifiability: In the case where requests come without the requester's identification, the provider might aim to associate an identifier to each of the requests. This could be the exact identifier of the actual requester; alternatively and less ambitiously, some assertion about the relationship between the request and the clients. As far as the provider succeeds in establishing a nontrivial and meaningful association, he would either learn precise personal data or obtain data that is somehow potentially personal and, thus, he would be able to compromise privacy to some extent.

Seen from the point of view of a client acting as requester, that client would be interested to estimate the extent of compromise achievable by the provider according to some suitable metric.

Location identification and classification: Whether by technical failures or by intentional blurring, a communicated location might differ from the actual one. Accordingly, the provider might aim to determine the actual location by some kind of reasoning, thereby strengthening his knowledge about the requester. As far as the requester is already identified, in this way the provider would obtain improved personal data regarding the requester. If the requester is so far not fully identified, more precise knowledge about the actual location might be helpful for the reidentification analysis or other analysis tasks.

Besides pure geographical data about the location, the provider might also aim to determine the kind of social activities offered at the respective place and thus learn information of the requester's activities, again potentially leading to even more crucial personal data, for instance if the provider can determine that the requests are coming from a cancer treatment center. Typically, social activities could be classified and denominated according to some ontology, e.g., distinguishing between shopping, medical care, entertainment, food services, sports, education, and so on, even possibly refined to subcategories and enhanced by further descriptive features.

Again, the client would like to evaluate the expected achievements, in particular in terms of the grade of success and the sensitivity of an identified location depending on its semantics, according to some metric.

Subtrajectory linkage: While strictly speaking a client only communicates location-time points, she or he actually provides information about her or his movements over the time, i.e., about the resulting trajectory. Accordingly, the provider might aim to reconstruct the actual trajectory in an approximative way in order to learn more about the client. As before, besides the pure geographical data about the full trajectory, he might additionally be interested in the semantics of the curve in terms of a suitable ontology that extends the ontology for single locations.

Reconstructing an actual trajectory necessarily includes linking single locations as communicated and subtrajectories obtained before as belonging to the same client. This need is clearly supported by already knowing the association of the requests with identifiers, but also conversely, if originally unknown, learning this association might be facilitated by having established links before.

Again, the client would like to evaluate the expected achievements, in particular in terms of the grade of success and the sensitivity of a reconstructed trajectory depending on its semantics.

Habitual Classification: Even if a set of requests is insufficient to link a meaningful subtrajectory or to clearly identify the user or the location, a provider might wish to classify the user according to certain metrics, e.g., for context-based advertising or to forecast the user behavior.

For example, it is possible to guess the social environment of a user by observing which ID's regularly issue requests from similar positions. A provider might find out the employment status by comparing the requests from workdays with the requests issued at weekends. It is also possible to learn about vacations, religion, etc.

The challenge remains to define and compile a complete list of the risks to privacy that Location-Based Services pose. Such a compilation will aid creation and analysis of privacy enhancing techniques designed to combat these risks by allowing better descriptions of the effect of these techniques.

A. Legal status of location privacy

The goal of introducing uncertainty in location when using LBS to help minimize violations of privacy can have legal implications as well, as location information can fall under the purview of privacy laws particularly dealing with identification or re-identification of data. An element in the definition of personal data in the EC Directive is that personal data indicates an identified or identifiable person. In other words, the terms "identified or identifiable" focus on the conditions under which an individual should be considered as "identifiable". In this regard the particular conditions of a specific case play an important role in this determination. Therefore the effect of uncertainty has to be addressed individually.

Location-based services in general process personal data in order to fulfill their contractual duties. The legal ground of using such information primarily is bound to the requirements of "informed consent" or "performance of a contractual duty under EC Directive 95/46 ([5], OJ L 281, p. 31 of 23.11.1995). Furthermore, processing on a secondary basis requires the fulfillment of at least one of the exceptions under the EC Directive such as the existence of a "legitimate interest" of the data processor or the existence of a "vital interest" of the data subject.

II. METHODS TO INDUCE UNCERTAINTY

It is reasonably straightforward for a mobile device to report something other than true location information. In fact, the Android operating system has a "debugging" feature allowing users to specify the exact longitude and latitude of the position returned to the application, although this feature does not extend to supplying the accuracy of this fake location. Applications exist on the Google Playstore to provide a friendly user interface for this feature [6], [7]. Software exists for the Android OS, when rooted, to expand on this capability and report user-specified information to applications rather than any actual information. This software extends to all forms of information, including the reported accuracy of a location.

Thus the question is not *can* we report an uncertain location, but how do we determine the uncertain location to report?

Work in this area attempts to decrease the certainty by which a location service can establish connections between specific users and their location data. In particular, we strive to degrade the confidence of these associations to the point that we can provide a well-defined level of protection against an adversary attempting to either (1) identify an individual from a particular location or a sequence of locations, (2) link a location or a sequence of locations to an individual, (3) connect a set of locations to a trajectory that belongs to the same (yet anonymous) individual, or (4) gain information about an individual based on habitual requests or locations. For that purpose, a wide number of anonymization approaches and obfuscation techniques exist. We subsume these techniques under the term “uncertainty methods”. The applicability of uncertainty methods depends on the kind of data that needs to be modified. Thus, we distinguish between the introduction of uncertainty to single locations and sequences of locations. While single locations are typical for one-shot queries sent to a location-based system, sequences of locations might be trajectories recorded by a smartphone with an activated GPS receiver or sequences of consecutive queries that have been sent from multiple positions to a location-based service.

The amount of uncertainty that is bound to a certain location or a trajectory depends on many factors. For example, if a pedestrian produces a location in the middle of a motorway or in a military exclusion area, an adversary might guess that this is implausible. Another example is a cyclist who has generated sequences of locations in distances that cannot be reached with the typical speed of a bicycle.

The actual effect of uncertainty may also be hindered by the apriori knowledge of the attacker rather than inherent attributes of the location itself. Beyond simple descriptions of location as a region of equal probability, an attacker likely possesses a number of maps of the uncertain location. Each of these maps imparts a form of semantics on the region that can suggest plausible positions within the uncertain location, thus reducing the true uncertainty achieved. An attacker is also able to potentially gain information about a group as a whole by viewing a collection of reported anonymous groups as the attacker gains knowledge of a large group of people in a relatively stable reporting area. The attacker is still unable to provide unique details about the target but are able to gain some insight about the group as a whole, the targeted individual included. This type of inference can still be a privacy violation even if the attacker is unable to define unique attributes. Assuming attackers have access to a history of location information, the attacker might be able to gain knowledge of similarity to previously identified points even if the new point is given with uncertainty. All of these factors can lead to information being revealed and privacy violated even if each single point is given with artificial uncertainty.

To ease our presentation, in the following we consider pairs of latitude, longitude only. However, all approaches described can be easily applied to more complex spatio-temporal settings by considering height and time as additional dimensions that are treated in the same way as latitude, longitude.

We do not presume to give a complete survey of work

in the area of location privacy, or a comprehensive listing of possible techniques, but rather a few choice examples to highlight the issues and possibilities.

A. Obfuscation techniques

Obfuscation techniques work to decrease the specificity of the location information provided to the service on behalf of a particular individual. In general, these methods can be applied by each user in isolation. One of the most intuitive techniques to increase the uncertainty of a location information is to add or multiply the latitude, longitude-record with a random number taken from a uniform distribution. The upper and lower bounds of the probability distribution function are a measure for the amount of uncertainty obtained.

A more sophisticated approach [8] takes the amount of uncertainty into account that has been already induced by the location sensor. In particular, the approach assumes that the correct position of a user is uniformly distributed over a circle that has been reported as center, radius by a GPS device. The radius specifies the uncertainty of the location information. This technique considers increases in privacy that can be gained by shifting the returned location or by modifying the radius of uncertainty rather than simply considering expansions of the radius.

Another method of obfuscation is the creation of a set of realistic dummies [9]. With this approach, the user not only sends a single position information to a location-based service, but a number of artificial dummy positions plus the real position. Accordingly, the service returns one result for each query. The client filters the set of results for the answer that corresponds to the real position. In this case, uncertainty is not defined as uncertainty towards a region (specified by a probability distribution), but as uncertainty towards which of the queried positions is the real one. Later work has focused on expanding the behavior of the generated dummies to better simulate the actual travel behavior of users to prevent the discovery of dummies. Behavior expansions include accounting for road networks over simple euclidean distance and travel constraints based on map information [10] and including pauses in the travel patterns of dummies [4]. As these techniques require that an adversary is unable to distinguish which queries are generated by dummies, such expansions of behavior add practical value to the privacy of these techniques.

Finally, there exist approaches to replace latitude, longitude-pairs with the positions of prominent landmarks [11]. For example, the exact position 49.530, 6.899 of a participant of a seminar in Dagstuhl could be reported as “Saarland”, “Germany” or “Europe”. The applicability of this kind of obfuscation depends on the format in which a location-based service requires location information.

B. Anonymization techniques

Anonymization techniques seek to obscure the relation between a user’s identity and his location information, by hiding his results amongst a set of other users.

Thus, anonymization requires the position information of multiple users. A popular approach is Spatio-Temporal Cloaking [12]. The approach adapts the concept of k -anonymity

for geographic coordinates. For this purpose, the approach computes cliques of users that are close together, and releases minimum bounding rectangles that contain the positions of at least k different users each. Various variants of this concept exist, e.g. peer-to-peer-anonymization [12]. These techniques can also be used for entire trajectories [13].

Mix zones [14] are an approach to add uncertainty to spatiotemporal settings where the users are continuously observed by a service provider. The approach identifies each user by a pseudonym. Furthermore, it divides regions into mix zones and application zones. In predefined time intervals, all individuals within a mix zone have the option to choose a new pseudonym. Given the number of users in a mix zone is large enough, the service provider cannot link the movement of an individual in one application zone to the movement of the same individual in another application zone.

The challenge is then to create a comprehensive set of obfuscation and anonymization techniques designed to limit the risk to privacy that can occur from secondary analysis of location data as highlighted in Section I.

III. CURRENT USE OF LOCATION UNCERTAINTY

The degradation in quality of service from intentionally giving uncertain locations is a critical issue and this is largely dependent on the particular application, and the implementation of that service provider. To evaluate the use of uncertainty, we first decompiled 19 applications on the Google Playstore and attempted to trace whether the reported uncertainty was used by the application and if so, how.

On a technical note, in the Android architecture there are two system calls, one that returns the location and a separate system call that returns the uncertainty of that location expressed in meters. Prior to this discussion there has been little impetus for the application developers to take into account any information aside from the reported location. Therefore, very few applications currently even make use of the extra information provided. This consideration will not concern us here, as these applications can simply be rewritten with the other system call in mind if needed. Our principle claim is that LBS can function in interesting ways in light of intentional uncertainty.

The applications that we considered generally fall into two categories: local search and check-in. Local search is a form of LBS application that does not greatly depend on accurate location so long as the uncertain data does not stray from the current location drastically. These types of applications do not generally ask for or make use of reported uncertainty. For instance, if a user's location is reported within a five-mile radius, the utility of Gas Buddy would not be greatly impaired. It would still display a list of gas stations that were all within fifteen miles of the user (the default for the application is ten miles, but can be lessened) regardless of the inaccurate data. Similarly, the local search functionality of the Yelp application would also not be greatly disrupted. More generally, it seems that most local search applications would not be rendered dysfunctional by intentional uncertainty.

A priori more likely to break are check-in applications, which use specific location information to ensure authenticity

TABLE I. ANDROID APPLICATIONS EVALUATED TO DETERMINE IF THEY MAKE A SYSTEM CALL TO DETERMINE UNCERTAINTY, AND IF THEY APPEAR TO SEND THAT INFORMATION TO THE SERVER.

Application	Syscall	Sends	Impacts Outcome	Local Search	Check-In
AroundMe				X	
Blendr					X
FourSquare	X	X	X		X
GasBuddy	X	X		X	
Google+					X
iExit				X	
Instagram					X
iRadar				X	
Layar					
Localicious	X			X	
Lookout	X				
Neer		X			
Poynt				X	
Radardroid				X	
SCVNGR					
Sygie					
TaxiMagic	X			X	
Yelp (check-in)	X	X	X		X
Yelp (local search)	X	X		X	
Zagat				X	

of the check-in. Indeed Foursquare, the most popular check-in application, in an effort to prevent “armchair mayors”, uses exact location information when possible, and applies a different set of criteria when the reported uncertainty is high [15], [16]. Therefore the functionality of this application would be greatly impaired by introducing artificial uncertainty. The Yelp application also has a check-in feature that uses a different standard for determining authenticity. The application computes the distance between your reported location and the check-in site, and compares it to the reported uncertainty [17]. Therefore, the Yelp application would continue to work exactly as before even with the new inputs. This would allow for check-ins while still maintaining privacy of the user. Indeed, it even allows for a new measure of privacy in that a user can timeshift their check-in by doing so before or after their visit, where the length of time is proportional to the intentional uncertainty introduced. Of course, this does pose a new problem in verifying the authenticity of such check-in, but nonetheless adds credence to the claim that LBS can function without exact location information.

Finally, it may be objected that some LBS, for instance turn-by-turn navigation and mapping applications, crucially rely on precise location information for their very utility, and that intentional uncertainty will render such services completely useless. While we hope to investigate solutions to this problem, we maintain that even if this is true, it does not impugn the validity or importance of the questions raised here. If we are successful in limiting the data released to other types of LBS, this still represents a significant gain from a privacy standpoint. The discussion will bear out exactly which LBS essentially require exact location data, and how to provide the others with uncertain data to protect the user's privacy.

The challenge lies in the real-world application of privacy enhancing techniques in current and future LBS applications while maintaining the quality of service that users expect from these Location-Based Services.

IV. CONCLUSIONS AND FUTURE WORK

We have seen that location data can be misused in potentially damaging ways: Location privacy is important. However,

location-based services are growing in popularity, and can provide significant value. The fact that location is inherently uncertain provides us an opportunity to maintain privacy while getting the benefit of location-based services. This opens opportunities for future research.

A first challenge is developing privacy definitions that provide practical protection. Different approaches may be needed depending on the privacy threat. For example, providing a distance-based region of uncertainty will protect against being found by a stalker; whereas a variable-distance uncertainty encompassing a k -anonymity style group may protect against a “Girls Around Me” style of application (assuming the adversary would not recognize which person you are within a group.) Given appropriate privacy definitions, we may find new challenges in implementing them, such as a mobile device independently (or perhaps in a peer-to-peer fashion) determining the uncertainty needed to encompass a k -anonymous group.

There is also the challenge of ensuring that applications make appropriate use of uncertainty data, ensuring good service despite privacy controls. Thankfully, applications do seem to recognize this data is available, and the fact that location may already be uncertain means that applications that use uncertainty information effectively are likely to win in the marketplace.

In summary, we believe that autonomous provision of privacy is both feasible, and can be done with limited impact on use of location-based services. This is an area where developments in privacy technology can have rapid transfer to real-world use, and provide meaningful privacy.

ACKNOWLEDGMENT

The authors wish to thank the participants of Dagstuhl Seminar 12331: Mobility Data Mining and Privacy for comments and discussions during the development of this work.

Nilgün Basalp, Bart Kuijpers, and Erkey Savas were supported by the MODAP (Mobility, Data Mining, and Privacy) project funded by EU, FP7, FET OPEN (www.modap.org).

Work of Joachim Biskup has been supported by the Deutsche Forschungsgemeinschaft (German Research Council) under grant SFB 876/A5 within the framework of the Collaborative Research Center “Providing Information by Resource-Constrained Data Analysis”.

Work of Chris Clifton supported by the NPRP grant 09-256-1-046 from the Qatar National Research Fund.

Shawn Merrill supported by the National Science Foundation under Grant No. 1012208.

The statements made herein are solely the responsibility of the authors.

REFERENCES

- [1] “Apple Q&A on location data,” Press Release, Apr. 27 2011. [Online]. Available: <http://www.apple.com/pr/library/2011/04/27Apple-Q-A-on-Location-Data.html>
- [2] J. Brownlee, “This creepy app isn’t just stalking women without their knowledge, it’s a wake-up call about facebook privacy [update],” *Cult of Mac*, Mar. 30 2012. [Online]. Available: <http://www.cultofmac.com/157641/this-creepy-app-isnt-just-stalking-women-without-their-knowledge-its-a-wake-up-call-about-facebook-privacy/>
- [3] E. Yigitoglu, M. L. Damiani, O. Abul, and C. Silvestri, “Privacy-preserving sharing of sensitive semantic locations under road-network constraints,” in *13th International Conference on Mobile Data Management*, 2012, pp. 186–195. [Online]. Available: <http://doi.ieeecomputersociety.org/10.1109/MDM.2012.48>
- [4] R. Kato, M. Iwata, T. Hara, A. Suzuki, X. Xie, Y. Arase, and S. Nishio, “A dummy-based anonymization method based on user trajectory with pauses,” in *SIGSPATIAL ’12 Proceedings of the 20th International Conference on Advances in Geographic Information Systems*, 2012, pp. 249–258. [Online]. Available: <http://doi.acm.org/10.1145/2424321.2424354>
- [5] “Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data,” *Official Journal of the European Communities*, vol. No I, no. 281, pp. 31–50, Oct. 24 1995. [Online]. Available: http://ec.europa.eu/justice_home/fsj/privacy/law/index_en.htm
- [6] M. Woo, “CatchMeIfUCan - fake location,” Android app, Jul. 12 2011. [Online]. Available: <https://play.google.com/store/apps/details?id=kr.woot0pia.gps>
- [7] Lexa, “Fake GPS location,” Android app, Sep. 20 2012. [Online]. Available: <https://play.google.com/store/apps/details?id=com.lexa.fakegps>
- [8] C. A. Ardagna, M. Cremonini, E. Damiani, S. D. C. di Vimercati, and P. Samarati, “Location privacy protection through obfuscation-based techniques,” in *Data and Applications Security XXI, 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security*, Redondo Beach, CA, USA, Jul. 8–11 2007, pp. 47–60. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-73538-0_4
- [9] 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 (ICDEW’05)*, Tokyo, Japan, Apr. 05–08 2005, p. 1248. [Online]. Available: <http://doi.ieeecomputersociety.org/10.1109/ICDE.2005.269>
- [10] A. Suzuki, M. Iwata, Y. Arase, T. Hara, X. Xie, and S. Nishio, “A user location anonymization method for location based services in a real environment,” in *GIS ’10 Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2010, pp. 398–401. [Online]. Available: <http://doi.acm.org/10.1145/1869790.1869846>
- [11] J. I. Hong and J. A. Landay, “An Architecture for Privacy-Sensitive Ubiquitous Computing,” in *Proceedings of the Second International Conference on Mobile Systems, Applications, and Services (MobiSys 2004)*, Jun. 06–09 2004. [Online]. Available: <https://dl.acm.org/citation.cfm?id=990087>
- [12] M. Gruteser and D. Grunwald, “Anonymous usage of location-based services through spatial and temporal cloaking,” in *Proceedings of First ACM/USENIX International Conference on Mobile Systems, Applications, and Services (MobiSys)*, May 2003.
- [13] M. E. Nergiz, M. Atzori, Y. Saygin, and B. Güç, “Towards trajectory anonymization: a generalization-based approach,” *Transactions on Data Privacy*, vol. 2, no. 1, pp. 47–75, 2009. [Online]. Available: <http://www.tdp.cat/issues/abs.a020a09.php>
- [14] A. R. Beresford and F. Stajano, “Mix Zones: User Privacy in Location-aware Services,” in *2nd IEEE Conference on Pervasive Computing and Communications Workshops (PerCom 2004 Workshops)*, Orlando, Florida, Mar. 14–17 2004, pp. 127–131. [Online]. Available: <http://doi.ieeecomputersociety.org/10.1109/PERCOMW.2004.1276918>
- [15] “On foursquare, cheating, and claiming mayorships from your couch...” blog, Apr. 7 2010. [Online]. Available: <http://blog.foursquare.com/2010/04/07/503822143/>
- [16] “The followup to our ‘mayorships from your couch’ post,” blog, Apr. 8 2010. [Online]. Available: <http://blog.foursquare.com/2010/04/08/505862083/>
- [17] J. B., “GPS vs Wifi: The battle for location accuracy using yelp check-ins,” blog, Aug. 14 2012. [Online]. Available: <http://engineeringblog.yelp.com/2012/08/gps-vs-wifi-the-battle-for-location-accuracy-using-yelp-check-ins.html>