22

Privacy-enhanced Location-based Access Control

C.A. Ardagna, M. Cremonini, S. De Capitani di Vimercati, P. Samarati

Dipartimento di Tecnologie dell'Informazione
Università degli Studi di Milano
26013 Crema, Italy
{ardagna,cremonini,decapita,samarati}@dti.unimi.it

Summary. Advancements in location technologies reliability and precision are fostering the development of location-based services that make use of the location information of users. An increasingly important category of such services is represented by Location-based Access Control (LBAC) systems that integrate traditional access control mechanisms with access conditions based on the physical position of users and other attributes related to the users location. Since privacy is extremely important for users, protection of their location information is paramount to the success of such emerging location-based services.

In this chapter, we first present an overview of Location-based Access Control systems and then characterize the location privacy protection problem. We then discuss the main techniques that have been proposed to protect location information, focusing on the obfuscation-based techniques. We conclude the chapter by showing a privacy-aware LBAC architecture and describing how a location-based access control policy can be evaluated.

1 Introduction

The widespread diffusion of pervasive technologies, as well as of mobile devices relying on them, makes available a great amount of high-sensitive location information that can be used for a variety of purposes. Customer-oriented applications, social networks and monitoring services can be functionally enriched with data reporting where people are, how they are moving or whether they are close by specific locations. To this end, several commercial and enterprise-oriented location-based services are already available and have gained popularity. Location-based services are supported by modern location technologies that have reached good precision and reliability at costs that most people (e.g., the cost of mobile devices) and companies (e.g., the cost of integrating location technologies in existing telecommunication infrastructures) can economically sustain. Since these location-based services are very complex and

may use the location information for different purposes, gathering and managing such information is a challenging aspect. Among the different issues that need to be addressed in the development of such services, *location privacy* is becoming increasingly important. Location privacy can be defined as the right of individuals to decide how, when, and for which purposes their location information could be released to other parties. The lack of location privacy protection could result in severe consequences that make users the target of fraudulent attacks such as [1]: *i) unsolicited advertising,* meaning that the location of the user could be exploited, without her consent, to provide advertisements of products and services available nearby the user position; *ii) physical attacks or harassment,* meaning that the location of the user could be used to carry physical assaults to individuals; *iii) users profiling,* meaning that the location of the user, which intrinsically carries personal information, could be used to infer other sensitive information such as state of health, personal habits, and professional duties; *iv) denial of services,* meaning that the location of the user could be used to deny accesses to services under some circumstances. In addition, location information can expose users to dangers such as stalking or physical harassment [2, 3].

Although location privacy is the subject of growing research efforts, there are no comprehensive solutions for location privacy protection in pervasive systems. The main branch of current research on location privacy focuses on users anonymity and on supporting online and mobile services that do not require the personal identification of a user for their provision [4, 5, 6]. When identification of users is required and, consequently, anonymity is not suitable, a viable solution to protect users privacy is to decrease the precision of personal information (including location) bound to identities [7, 8, 9]. For several online services personal information associated with identities does not need to be as accurate as possible to guarantee a certain service quality.

In this chapter, the issue of protecting location privacy is analyzed in the context of Location-based Access Control (LBAC) systems [10]. The remainder of this chapter is organized as follows. Section 2 presents basic concepts behind location-based access control systems. Section 3 provides a brief overview of different types of location privacy that must be preserved depending on the scenarios and on the requirements together with a description of the techniques that can be used to protect location privacy. Section 4 describes some obfuscation-based techniques aimed at privacy protection. Section 5 presents a privacy-aware LBAC architecture and discusses how the evaluation of location-based predicates can be performed. Finally, Section 6 gives our conclusions.

2 Location-based Access Control Systems

Novel access control mechanisms are based on the assumption that properties characterizing a requester, which are usually provided through *digital cre-*

dentials, are sufficient to decide which actions the requester is authorized to perform on resources [11]. However, requester's credentials are not the only information that should be considered in access control decisions. The rapid development in the field of wireless and mobile networking fostered a new generation of devices suitable for being used as sensors by location technologies, which are able to compute the relative position and movement of users in their environment. Therefore, the location of users, potentially available to access control modules, may also play an important role in determining access rights and allows the definition of a new class of location-based policies regulating access to and fruition of resources. When evaluating location-based access control policies, however, we need to consider that location-based information presents some peculiarities: location information is both *approximate* (all location system have a margin of error) and *time-variant* (the user position changes over time due to the on-going motion of requesters).

Location-based Access Control (LBAC) systems provide the infrastructure for managing and evaluating access control policies that include predicates and conditions based on the location information of users. LBAC systems should be designed to tolerate rapid context changes, because users are no longer forced to be at pre-defined fixed positions but they can freely access services through their mobile devices (e.g., mobile phones).

2.1 Location-based Conditions and Predicates

The first step towards the development of a LBAC system consists in the definition of location-based conditions. We identify three main classes of location-based conditions, which might be useful to include in access control policies and whose evaluation is possible with today's technology [10]:

- *position-based* conditions on the location of the user (e.g., to evaluate whether a user is in a certain building or city or in the proximity of other entities);
- *movement-based* conditions on the mobility of the users (e.g., velocity, acceleration, or direction where users are headed);
- *interaction-based* conditions relating multiple users or entities (e.g., the number of users within a given area).

The language presented in [10] supports such conditions and is based on the assumption that each user, who is unknown to the service responsible for location measurements, is univocally identified via a user identifier (UID). A unique identifier is also associated with physical and/or moving entities that may need to be located (e.g., a vehicle with an on-board GPRS card). A typical UID for location-based applications is the SIM number linking the user's identity to a mobile terminal. Moreover, the language is also based on the assumption that there is a set of map regions identified either via a geometric model (i.e., a range in a n-dimensional coordinate space) or a

Table 1. Examples of location-based predicates

Type	Predicate	Description
Position	inarea(*user, area*)	Evaluate whether *user* is located within *area*.
	disjoint(*user*, *area*)	Evaluate whether *user* is outside *area*.
	distance(*user*, *entity*, *min_dist*, *max_dist*)	Evaluate whether distance between *user* and *entity* is within interval [*min_dist*, *max_dist*].
Movement	velocity(*user*, *min_vel*, *max_vel*)	Evaluate whether *user*'s speed falls within range [*min_vel*, *max_vel*].
Interaction	density(*area*, *min_num*, *max_num*)	Evaluate whether the number of users currently in *area* falls within interval [*min_num*, *max_num*].
	local_density(*user*, *area*, *min_num*, *max_num*)	Evaluate the density within a 'relative' area surrounding *user*.

symbolic model (i.e., with reference to entities of the real world such as, for example, cells, streets, cities, zip code or buildings) [12].

Predicates are expressed as boolean queries of the form *predicate*(*parameters, value*). Table 1 illustrates some examples of location predicates.

Example 1. Let `alice` be a user identifier, and `Milan` and `Director_Office` be two map regions. Three simple examples of location-based conditions are the following.

- inarea(`alice`,`Milan`): request `alice` to be located in `Milan`.
- velocity(`alice`,70,90): request `alice` to travel at a speed included in the interval [70,90].
- density(`Director_Office`,0,1): request at most one person in the `Director_Office`.

2.2 Location-based Access Control Policies

Location-based access control policies can be considered as a means for enriching the expressive power of existing access control languages (e.g., [11, 13, 14, 15]) by introducing location-based predicates. We assume access control rules to be triples whose elements are generic boolean formula over the subject, object, and action domains. Formally, an access control rule is defined as follows.

Definition 1 (Access control rule). *An access control rule is a triple of the form* ⟨subj_expr, obj_expr, action⟩*, where:*

- subj_expr *is a boolean formula of terms referring to a set of subjects depending on whether they satisfy or not certain conditions that can evaluate the user's profile/information, location predicates, or the user's membership in groups, active roles, and so on;*

Table 2. Examples of access control rules regulating access to the Mobile Network Console and databases

	subject		action	object
	generic conditions	location-based conditions		
1	**user**.role=admin ∧ valid(**user**.username, **user**.password)	inarea(**user**.sim, Server_Room) ∧ density(Server_Room, 1, 1) ∧ velocity(**user**.sim, 0, 3)	execute	**object**.name=MNC
2	**user**.role=admin ∧ valid(**user**.username, **user**.password)	inarea(**user**.sim, Inf._System_Dept.) ∧ local_density(**user**.sim, Close_By, 1, 1) ∧ velocity(**user**.sim, 0, 3)	read	**object**.category= Log&Bill
3	**user**.role=CEO ∧ valid(**user**.username, **user**.password)	local_density(**user**.sim, Close_By, 1, 1) ∧ inarea(**user**.sim, Corp._Main_Office) ∧ velocity(**user**.sim, 0, 3)	read	**object**.category= customer
4	**user**.role=CEO ∧ valid(**user**.username, **user**.password)	local_density(**user**.sim, Close_By, 1, 1) ∧ disjoint(**user**.sim, Competitor_Location)	read	**object**.category= StatData
5	**user**.role=guest ∧ valid(**user**.username, **user**.password)	local_density(**user**.sim, Close_By, 1, 1) ∧ inarea(**user**.sim, Corporate_Location)	read	**object**.category= StatData

- obj_expr *is a boolean formula of terms referring to a set of objects depending on whether they satisfy or not certain conditions that can evaluate membership of the object in categories, values of properties on metadata, and so on;*
- action *is the action (or class of actions) to which the rule refers.*

Each profile is referenced with the identity of the corresponding user/object. Single properties within users and objects profiles are referenced with the traditional dot notation. For instance, `alice.address` indicates the address of user `alice`. Here, `alice` is the identity of the user (and therefore the identifier for the corresponding profile), and `address` is the name of the property. To refer to the user and the object involved in a request without introducing variables in the language, we use two keywords: **user** indicates the identifier of the person making the request; **object** indicates the identifier of the object to which access is requested.

Example 2. Consider a company responsible for the management of a mobile network that needs both strong authentication methods and expressive access control policies. Suppose that the Mobile Network Console (MNC) is the software that permits to reconfigure the mobile network. Managing a nationwide mobile network is an extremely critical activity because reconfiguration privileges must be granted to strictly selected personnel only and must be performed according to high security standards (rule 1 in Table 2). In addition to reconfiguration privileges, also the access to mobile network's databases must be managed carefully and according to different security standards depending on the level of risk of the data to be accessed. In particular, access to logging and billing data is critical, because they include information about the position and movements of mobile operator's customers (rule 2 in Table 2). Access to customer-related information is usually less critical but still to be

handled in a highly secured environment and to be granted only to selected personnel, according to the laws and regulations in force (rule 3 in Table 2). Finally, access to statistical data about the network's operation is at a lower criticality level, whereas they are still private information to be protected, for example, from disclosure to competitors (rules 4 and 5 in Table 2).

In the following, we discuss location privacy issues and present a location privacy solution suitable for location-based services along with a privacy-aware LBAC architecture.

3 Location Privacy

Although location information can be exploited for providing enhanced services, the high sensitivity of such an information increases concerns of users about their privacy. Location privacy can assume several meanings and pursue different objectives, depending on the services the users are interacting with. The following categories of location privacy have been identified.

- *Identity privacy.* The main goal is to protect users' identities that could be directly or indirectly inferred from location information [4, 5, 6, 16]. To this purpose, protection techniques aim at minimizing the disclosure of the data that can let an attacker infer a user identity, such as home and work addresses. This type of location privacy is suitable in application contexts that do not require the identification of the users as a fundamental information for service provisioning. For instance, many online services provide a person with the ability to establish a relationship with some other entities (e.g., anonymous chats) or with some applications (e.g., allergy warning) without her personal identity being disclosed to that entity. In this case, the best possible location measurement can be provided to the others entities but the actual user's identity must be preserved.
- *Position privacy.* The main goal is to protect the position information of individual users, by perturbing corresponding information and decreasing the accuracy of location information [7, 8, 9]. Position privacy is suitable for environments where users' identities are required for a successful service provisioning, and less accurate location information does not severely affect the service quality (e.g., access to services inside a production plant or friends finder services). A technique that most solutions exploit, either explicitly or implicitly, consists in reducing the accuracy by scaling a location to a coarser granularity (e.g., from meters to hundreds of meters, from a city block to the whole town).
- *Path privacy.* The main goal is to protect the privacy of information associated with users motion, such as the path followed while traveling or walking in a urban area [17, 18, 19]. There are several location-based services (e.g., personal navigation systems) that could be exploited to subvert path privacy or to illicitly track users. Path privacy is the most complex

class of location privacy problem and can refer to identity privacy and/or position privacy.

The above three privacy categories pose different requirements that are fulfilled by different techniques. The heterogeneity of location privacy problems results then in a lack of a general solution able to satisfy all the privacy requirements. In the following, different classes of techniques are discussed and analyzed.

3.1 Location Privacy Techniques

Location privacy techniques can be partitioned into three main classes that correspond to the different types of location privacy above-mentioned: *anonymity-based*, *policy-based*, and *obfuscation-based*. These classes are partially overlapped in scope and could be potentially suitable to cover requirements coming from one or more of the categories of location privacy. Anonymity-based and obfuscation-based techniques can be usually regarded as dual categories. While anonymity-based techniques have been primarily defined to protect identity privacy and are less suitable for protecting position privacy, obfuscation-based techniques are well suited for position protection and less appropriate for identity protection. Anonymity-based and obfuscation-based techniques are well-suited for protecting path privacy. Nevertheless, more studies and proposals have been focused on anonymity-based rather than on obfuscation-based techniques. Policy-based techniques are in general suitable for all the location privacy categories; however, they can be difficult to understand and manage for end users.

Anonymity-based techniques

This class of techniques focus both on identity privacy and path privacy protection [4, 5, 6, 20]. Beresford and Stajano [4, 21] propose a *mix zone* model and employs an anonymity service based on an infrastructure that delays and reorders messages from subscribers within pre-defined zones. The mix zone model is based on a trusted middleware positioned between location systems and third party applications, which is responsible for limiting the information collected by applications. An application selects a set of *application zones* representing application interests in specific geographic areas, such as hospital, supermarket, and so on. Users register interest in a specific set of applications and the middleware limits the location information that such applications can receive to the locations inside the application zones. Each user has one or more unregistered geographical regions, called *mix zones*, where users cannot be tracked, that is, when a user enters a mix zone her identity is mixed with all other users in the same mix zone. The mix zones model is then aimed at protecting long-term user movements still allowing the interaction with many location-based services. However, the effectiveness of such a solution is

strongly dependent on the number of users joining the anonymity service and, in particular, on the number of users physically co-located in the same mix zone at the same time.

Bettini et. al. [5] propose a framework able to evaluate the risk of sensitive location-based information dissemination and introduce a technique aimed at supporting k-anonymity [8, 9]. The concept of k-anonymity captures a traditional requirement of statistical agencies stating that released data must be indistinguishably related to no less than a certain number (k) of users. Traditionally, k-anonymity is based on the definition of a *quasi-identifier* that is a set of attributes exploitable for linking data to identifiers. The k-anonymity requirement states that each release of data must guarantee that every combination of values of quasi-identifiers can be indistinctly linkable to at least k individuals. The proposal in [5] puts forward the idea that the geo-localized history of the requests submitted by a user can be considered as a quasi-identifier that can be used to discover sensitive information about that user. For instance, a user tracked during working days is likely to commute from her house to the workplace in a specific time frame in the morning and to come back in another specific time frame in the evening. This information could be used to identify the user. Consequently, the service provider gathering both user requests for services and personal history of locations (i.e., a sequence of user location updates) should never be able to link a subset of requests to a single user. To make this possible, there must exist k users having a personal history of locations consistent with the set of requests that have been issued. This solution is highly dependent on the availability of k indistinguishable histories of locations: the worst case happens when a given user has a unique history, which make her always identifiable.

Also other proposals [6, 20] rely on the concept of k-anonymity by requiring that a user should be indistinguishable from other $k-1$ users in a given spatial area or temporal interval. Gruteser and Grunwald [6] propose a middleware architecture and an adaptive algorithm to adjust location information resolution, in spatial or temporal dimensions, to comply with specified anonymity requirements. To this purpose, the authors introduce the concepts of *spatial* and *temporal cloaking* used to transform the location of a user to a different location that satisfies the required level of anonymity. Spatial cloaking guarantees k-anonymity by applying an adaptive quad-tree algorithm that decreases the spatial resolution to an area that contains k indistinguishable users. Temporal cloaking, which is orthogonal to the spatial cloaking, provides spatial coordinates with higher accuracy but it reduces the accuracy in time. The key feature of the adaptive cloaking algorithm is that the required level of anonymity can be achieved for any location. Mokbel et al. [20] present a framework, named *Casper*, that changes traditional location-based servers and query processors to provide the users with anonymous services. Users can define their privacy preferences through two parameters: k, meaning that the user wants to be indistinguishable among other k entities; and A_{min} representing the minimal area that the user is willing to release. The core of the

Casper framework is composed by two components: a *location anonymizer,* which is responsible for perturbing the user location until user's privacy preferences are satisfied, and a *privacy-aware query processor,* which is responsible for the management of anonymous queries and cloaked spatial areas.

Anonymity-based techniques have also been exploited to guarantee *path privacy* protection [17, 18, 19]. In particular, path privacy involves the protection of users that are in motion and are continuously monitored during a time interval. This research field is particularly relevant for location tracking applications designed and developed for devices with limited capabilities (e.g., cellular phones), where data about users moving in a particular area are collected by external services. Gruteser et al. [17] propose a solution to path privacy protection by means of *path anonymization functions*. The authors argue that the association of a single or multiple pseudonyms, which change over time, with a user is not sufficient to provide path privacy protection. Privacy provided by pseudonyms can be actually subverted by applying an inference process that gathers path information, such as the place a user stays during the night. Therefore, since it is difficult to provide strong anonymity for path protection because it would require the existence of several users traveling along the same path at the same time, Gruteser et al. provide two techniques that guarantee a "weaker anonymity", meaning that users could potentially be linked to their identities but at price of huge computational efforts. The first technique relies on *path segmentation,* which partitions a user's path in a set of smaller paths changing, at the same time, the associated pseudonym. The second technique relies on *minutiae suppression* that suppresses those parts of a path that are more distinctive and could bring to an easy association between a path and an identity. The suitability of these techniques is highly dependent on the density of users in the area in which the adversary collects location samples. In areas with low density of users, an adversary has a good likelihood of tracking individuals, whereas in areas with many overlapping paths, linking segments to identities can be extremely difficult.

Other proposals consider path protection as a process whose outcome must be managed by a service provider and consequently privacy techniques have to preserve a given level of accuracy to permit a good quality of service provisioning. Gruteser and Liu [18] present a solution based on the definition of a *sensitivity map* composed by sensitive and insensitive zones. Sensitive zones are those area where the users prefer to hide their visits. The work defines three algorithms aimed at path privacy protection: *base*, *bounded-rate*, and k-area. Among the three, the k-area algorithm stands out, giving the best performance in terms of privacy, and minimizing the number of location updates suppression. In particular, the k-area algorithm is built on top of sensitivity maps that are composed of areas containing k sensitive zones. Location updates of a user entering a region with k sensitive areas are temporarily stored and not released. If a user leaving that region has visited at least one of the k sensitive areas, location updates are suppressed; they are

released, otherwise. Finally, Ho and Gruteser [19] propose a path confusion algorithm. This algorithm introduces a level of uncertainty by creating cross paths between at least two users. In this case, the attacker observing different paths is not able to recognize which path has followed one specific user.

Policy-based techniques

Another class of location privacy techniques relies on the definition of *privacy policies*. Privacy policies define restrictions that regulate the release of the location of a user to third parties. Hauser and Kabatnik [22] address the location privacy problem in a privacy-aware architecture for a global location service, which allows users to define rules that will be evaluated to regulate access to location information. The IETF Geopriv working group addresses privacy and security issues related to the disclosure of location information over the Internet [23]. The main goal of the Geopriv working group is to define an environment (i.e., architecture, protocols, and policies) supporting both location information and policy data. Others works [24, 25] used the Platform for Privacy Preferences (P3P) [26] to encode users privacy preferences.

In summary, policy-based techniques allow a flexible definition of policies that fit the user needs of privacy by restricting the ability to manage locations and disclosing information. However, although policies-based solutions are suitable for privacy protection, users are often not willing to directly manage complex policies and, hence, may refuse participation in pervasive environments.

Obfuscation-based techniques

Obfuscation-based techniques are aimed at protecting location privacy by degrading the accuracy of the location information still maintaining an explicit association with the real user identity.

Duckham and Kulik [7] define a framework that provides a mechanism for balancing individual needs for high-quality information services and location privacy. The proposed solution is based on the concept of *imprecision*, which indicates the lack of specificity of location information. The authors suggest to degrade location information quality and to provide obfuscation features by adding n points with same probability of being the real user position. The algorithm assumes a graph-based representation of the environment. Also, the authors propose a validation and evaluation of their methods through a set of simulations [27]. The results show that obfuscation can provide at the same time a high service quality and a high privacy level.

Other proposals relies on a trusted middleware, which lies between location providers and location-based applications, responsible for enforcing users privacy preferences before releasing location information. Openwave [28], for example, includes a location gateway that obtains users location information

from multiple sources and delivers them, possibly modified according to privacy requirements, to other parties. Users define their privacy preferences in terms of a minimum distance representing the maximum location accuracy they are willing to accept. Bellavista et al. [29] present a solution based on a middleware that balances the level of privacy requested by users and the need of service precision. Location information is perturbed depending on privacy/efficiency requirements negotiated by the parties and it is returned with lower precision and lower geographical granularity.

In summary, although obfuscation-based techniques are compatible with users specifying their privacy preferences in a common and intuitive manner (usually as a *minimum distance*), they do not provide a quantitative estimation of the provided privacy level, and they usually implement a single obfuscation technique, which provide an obfuscation effect by scaling up the extent of the location area.

4 Obfuscation Techniques for Location Privacy Protection

An interesting research direction is to use obfuscation-based techniques for location privacy protection in LBAC systems [30, 31, 32]. These recent proposals provide privacy by degrading the location accuracy of each measurement while offering a measurable accuracy to service providers and are based on two working assumptions that simplify the analysis with no loss of generality: *i)* the area returned by a location measurement is planar and circular, which is the actual shape resulting from many location technologies; *ii)* the distribution of measurement errors within a returned area is uniform. The first assumption derives from the fact that user location information is affected by an intrinsic measurement error introduced by sensing technologies, resulting in spatial areas rather than geographical points. This assumption represents a particular case of the general requirement of considering convex areas and a good approximation for actual shapes resulting from many location technologies (e.g., cellular phones location). A location measurement is then defined as follows.

Definition 2 (Location measurement). *A location measurement of a user u is a circular area $Area(r, x_c, y_c)$, centered on the geographical coordinates (x_c, y_c) and with radius r, which includes the real user's position (x_u, y_u) with probability $P((x_u, y_u) \in Area(r, x_c, y_c)) = 1$.*

Definition 2 comes from observing that sensing technologies based on cellular phones usually guarantee that the real user's position falls within the returned area.

The second assumption is introduced to discuss the effects of obfuscation techniques. Consider a random location within a location measurement

$Area(r, x_c, y_c)$, where a "random location" is a neighborhood of random point $(\hat{x}, \hat{y}) \in Area(r, x_c, y_c)$. The probability that the real user's position (x_u, y_u) belongs to a neighborhood of a random point $(\hat{x}, \hat{y})$ is uniformly distributed over the whole location measurement. Accordingly, the joint probability density function (pdf) of the real user's position can be defined as follows.

Definition 3 (Uniform joint pdf). *Given a location measurement $Area(r, x_c, y_c)$, the joint probability density function (joint pdf) $f_r(x, y)$ of real user's position (x_u, y_u) to be in the neighborhood of point (x, y) is:*

$$f_r(x, y) = \begin{cases} \frac{1}{\pi r^2} & \text{if } (x, y) \in Area(r, x_c, y_c) \\ 0 & \text{otherwise.} \end{cases}$$

Before analyzing the obfuscation techniques in details, we first describe how users can express their privacy preferences. Despite its importance for the effectiveness of a privacy solution, this issue has received little attention in previous works on location privacy. We then describe how the level of privacy can be quantitatively expressed as a functional term independently from any physical scale or specific technology.

4.1 User Preferences and Relevance Metric

Several works in location privacy field are based on the definition of users privacy preferences by means of a minimum distance [7, 28]. This choice is dictated by the fact that usually the users tend to adopt simple and intuitive way for expressing their privacy preference and tend to be averse to complex configurations. A user can define as her privacy preference a minimum distance, which results in a location area achieved by increasing the granularity of the actual location measurement. In particular, assuming location measurements as circular areas, the minimum distance privacy preference represents the minimum radius of the area that a user is willing to release to other parties. However, the definition of the minimum distance as user privacy preference exhibits some shortcomings: *i)* it is highly dependent on the adopted privacy solution; *ii)* it is suitable for only obfuscation techniques that increase the granularity of the measurement; *iii)* it is difficult to integrate in a full-fledged location-based application scenario [10, 33]; *iv)* it is not suitable for solutions using different obfuscation techniques.

To overcome these issues, others proposals [30, 31, 32] suggest a different way to manage users privacy preferences. In these works, users specify their privacy requirements through the definition of a *relative* degradation of the location accuracy with respect to the location measurement, which is modeled through an index $\lambda \in [0, \infty)$, where $\lambda = 0$ corresponds to no degradation, $\lambda \to \infty$ to maximum degradation, and intermediate values correspond to different degrees of degradation. For instance, λ=0.5 means 50% of degradation,

λ=1 means 100% of degradation and any value λ >1 corresponds to a degradation greater than 100%. Although both minimum distance d and index λ are easy to specify for users, λ is a more general solution because independent from a specific location measurement and obfuscation technique. However, the definition of λ is not sufficient, especially when we need to balance the users needs of privacy and the LBSs needs of location accuracy to maintain an acceptable quality of the online service.

To accommodate the peculiar characteristics of the above scenario, the concept of *relevance* is introduced as the adimensional metric of both the accuracy and the privacy of a location information, abstracting from any physical attribute of sensing technology. A relevance $\mathcal{R}$ is a value in (0,1] associated with each location information, which depends on measurement errors and privacy preferences of users. In particular, $\mathcal{R}$ tends to 0 when the location information is considered unreliable for service provision; $\mathcal{R}$=1 when the location information is equal to the original location measurement; $\mathcal{R} \in$ (0,1) when the location information has various degrees of accurateness. The location privacy associated with an obfuscated location is evaluated by (1-$\mathcal{R}$).

Applying the concept of relevance to a LBAC scenario, an LBAC service has to manage the following different relevances:

- *Technological relevance* ($\mathcal{R}_{Tech}$) is the metric for the accuracy of the location measurement provided by a location service given a mobile technology and its technical quality.
- *Privacy relevance* ($\mathcal{R}_{Priv}$) is the metric for the accuracy of an obfuscated location and therefore the level of privacy provided to the users.
- *LBAC relevance* ($\mathcal{R}_{LBAC}$) is the metric for the lowest accuracy of the location information that an LBAC service is willing to accept. It is required by the business application for a location measurement or for a location-based predicate evaluation.
- *Evaluation relevance* ($\mathcal{R}_{Eval}$) is the metric for the accuracy of a LBAC predicate evaluation.

Among these relevances, $\mathcal{R}_{LBAC}$ and $\mathcal{R}_{Tech}$ are assumed to be known. $\mathcal{R}_{Priv}$ is derived from the privacy preferences expressed by users, while $\mathcal{R}_{Eval}$ is calculated by the system (see Sect. 5). In other words, $\mathcal{R}_{Priv}$ represents the relevance of the final obfuscated area that is calculated starting from the location measurement with relevance $\mathcal{R}_{Tech}$ and by degrading its accuracy according to the value of λ. Formally, $\mathcal{R}_{Priv}$ is calculated as:

$$\mathcal{R}_{Priv} = (\lambda + 1)^{-1} \mathcal{R}_{Tech} \tag{1}$$

If a privacy preference is expressed through a minimum distance r, it is straightforward to derive λ from r. The obfuscated area is then calculated by scaling up the radius of the location measurement until the user privacy preference λ is satisfied.

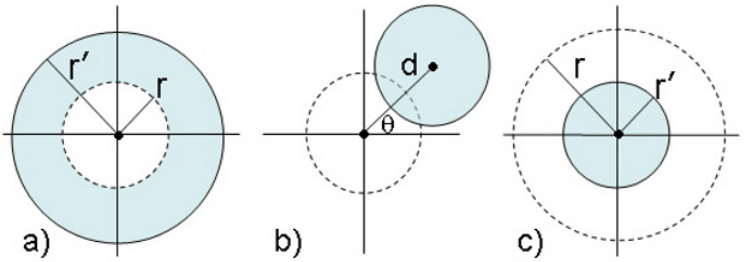

Fig. 1. Enlarging (a), shifting (b), and reducing (c)

4.2 Obfuscation Techniques

We present some obfuscation techniques that reduce the location accuracy of a location measurement until the privacy preferences are achieved. In particular, each technique takes λ as input and computes $\mathcal{R}_{Priv}$ and the obfuscated area.

Enlarging the Radius

Enlarging the radius of a location measurement represents the traditional solution adopted in the context of location privacy protection. Given a location measurement $Area(r, x_c, y_c)$, an obfuscated area $Area(r', x_c, y_c)$ is generated, where $r' > r$ (see Fig. 1(a)). The obfuscation effect directly derives from the fact that the joint pdf associated with the obfuscated area decreases, that is, $\forall r, r' \in \mathbb{R}^+ : r < r' \Rightarrow f_r(x, y) > f_{r'}(x, y)$. The relevance $\mathcal{R}_{Priv}$ of the location information after spatial obfuscation can be derived from $\mathcal{R}_{Tech}$ by considering the ratio of the two pdf as a scalar factor:

$$\mathcal{R}_{Priv} = \frac{f_{r'}(x, y)}{f_r(x, y)} \cdot \mathcal{R}_{Tech} = \frac{r^2}{r'^2} \cdot \mathcal{R}_{Tech}, \quad \text{with } r < r' \tag{2}$$

Given a privacy preference $\lambda \geq 0$, the radius of the obfuscated area r' is calculated from (1) and (2) as follows:

$$r' = r\sqrt{\lambda + 1}$$

This relation permits to generate the obfuscated area by enlarging radius r to radius r', which satisfies, according to our semantics, the user privacy preference λ. Note that, if the privacy preference of the user is provided by means of a minimum distance (i.e., radius r') relevance $\mathcal{R}_{Priv}$ of the obfuscated area is always calculated by equation (2).

Shifting the Center

Shifting the center of the area returned by a sensing technology is another way of obfuscating a location measurement. The obfuscated area is derived from the original area by calculating the distance d between the two centers and the shifting angle θ. Let $Area(r, x_c + \Delta x, y_c + \Delta y)$ be the obfuscated area. Note that, since LBAC applications cannot deal with false information to provide a service, obfuscated areas with no intersection with the original location measurement are considered not acceptable. The reason is that, since location measurements contain users positions with probability 1, all the areas disjoint with a location measurement have probability 0 of including the real user location, and then are indiscernible using the relevance metric. Therefore, these areas must be simply considered as false location information.

The privacy gain can be measured by considering the intersection of the original and obfuscated areas, denoted $Area_{Tech \cap Priv}$. Intuitively, the degree of privacy is inversely proportional to the intersection of the two areas and therefore it is directly proportional to the distance $d \in [0, 2r]$ between the two centers. In particular, if $d = 0$, there is no privacy gain and $P((x_u, y_u) \in Area(r, x_c + \Delta x, y_c + \Delta y)) = P((x_u, y_u) \in Area(r, x_c, y_c)) = 1$. If $d = 2r$, there is maximum privacy and $P((x_u, y_u) \in Area(r, x_c + \Delta x, y_c + \Delta y))$ tends to 0; and if $0 < d < 2r$, there is an increment of privacy and $0 < P((x_u, y_u) \in Area(r, x_c + \Delta x, y_c + \Delta y)) < 1$.

Angle θ (see Fig. 1(b)) is assumed to be randomly chosen, since all values of θ are equivalent with respect to the privacy preferences of users.

To measure the obfuscation effect and define the relation between relevances, two probabilities must be composed: *i)* the probability that the real user's position belongs to the intersection $Area_{Tech \cap Priv}$, and *ii)* the probability that a random point selected from the whole obfuscated area belongs to the intersection. Then, the relation between relevances $\mathcal{R}_{Tech}$ and $\mathcal{R}_{Priv}$ is represented by:

$$\mathcal{R}_{Priv} = P((x_u, y_u) \in Area_{Tech \cap Priv}) \cdot P((x, y) \in Area_{Tech \cap Priv}) \cdot \mathcal{R}_{Tech} = \frac{Area_{Tech \cap Priv}}{Area(r, x_c, y_c)} \cdot \frac{Area_{Tech \cap Priv}}{Area(r, x_c + \Delta x, y_c + \Delta y)} \cdot \mathcal{R}_{Tech} = \frac{Area^2_{Tech \cap Priv}}{Area(r, x_c, y_c)^2} \cdot \mathcal{R}_{Tech} \quad (3)$$

Given the privacy preference expressed by $\lambda \geq 0$, the distance d between the centers of the original and obfuscated area is calculated from (1) and (3) as follows:

$$(\lambda + 1)^{-1} = \frac{Area^2_{Tech \cap Priv}}{Area(r, x_c, y_c)^2} \quad (4)$$

The distance d between the centers is the unknown variable to be derived to obtain the obfuscated area. It can be calculated by expanding the term $Area_{Tech \cap Priv}$ as a function of d and by solving the following system of equations, whose variables are d, σ and γ. σ and γ are the central angles of circular

sectors identified by the two radii connecting the centers of the areas with the intersection points of original and obfuscated areas.[1]

$$\begin{cases} \left[\frac{\sigma}{2}r^2 - \frac{r^2}{2}\sin\sigma\right] + \left[\frac{\gamma}{2}R^2 - \frac{R^2}{2}\sin\gamma\right] = \sqrt{\delta}\pi r \cdot R \\ d = r\cos\frac{\sigma}{2} + R\cos\frac{\gamma}{2} \\ r\sin\frac{\sigma}{2} = R\sin\frac{\gamma}{2} \end{cases} \tag{5}$$

Solutions of this system can be obtained numerically.

Reducing the Radius

The third obfuscation technique consists in reducing the radius of a location measurement from r to r', as showed in Fig. 1(c). The obfuscation effect is produced by a correspondent reduction of the probability to find the real user location within the returned area, whereas the joint pdf is fixed.

Let (x_u, y_u) be the real user position coordinates, By assumption, the probability that the real user position falls in the location measurement of radius r is $P((x_u, y_u) \in Area(r, x, y)) = 1$. When we obfuscate by reducing the radius, an area of radius $r' < r$ is returned, where $P((x_u, y_u) \in Area(r', x, y)) < P((x_u, y_u) \in Area(r, x, y))$, since a circular ring having pdf greater than zero has been excluded.

With regard to relevances $\mathcal{R}_{Tech}$ and $\mathcal{R}_{Priv}$, their relation can be defined as:

$$\mathcal{R}_{Priv} = \frac{P((x_u, y_u) \in Area(r', x, y))}{P((x_u, y_u) \in Area(r, x, y))} \cdot \mathcal{R}_{Tech} = \frac{r'^2}{r^2} \cdot \mathcal{R}_{Tech}, \quad \text{with } r' < r \tag{6}$$

Given a privacy preference $\lambda \geq 0$, the radius of the obfuscated area r' is calculated from (1) and (6) as follows:

$$r' = \frac{r}{\sqrt{\lambda + 1}}$$

This relation permits to generate the obfuscated area by reducing radius r to radius r', which satisfies, according to our semantics, the user privacy preference λ.

5 Integrating Obfuscation Techniques with LBAC Systems

The definition of LBAC systems poses some architectural and functional issues that were never studied before in the context of traditional access control

[1] The system of equation (5) is presented in the most general form of two areas with different radii (i.e., r and R).

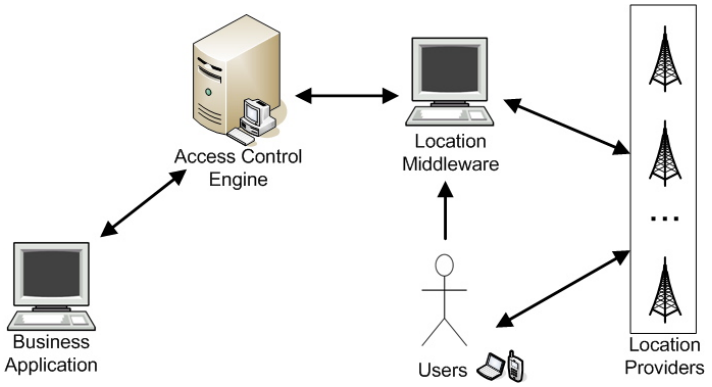

Fig. 2. A privacy-Aware LBAC Architecture

systems. A privacy-aware LBAC architecture must be developed integrating components logically tied with the applications that need location-based access control enforcement and components providing privacy-aware location services. One typical approach in the design of LBAC architectures is to provide a location middleware acting as a trusted gateway between the LBAC system and the location services. Such a component is in charge of managing all interactions with sensing technologies and enforce users privacy preferences. In [30, 31, 32] the authors present a privacy-aware LBAC architecture (see Fig. 2) whose logical components can be summarized as follows.

- *User.* It is the subject to be located through her mobile device during the interaction with the Business Application. The user first defines her privacy preferences at the Location Middleware and then interacts with the Access Control Engine to gain the access to the Business Application.
- *Business application.* It represents a service provider that offers resources protected by LBAC policies. It relies on the Access Control Engine for evaluating policies based on users location.
- *Access Control Engine (ACE).* It is the component responsible for the evaluation and enforcement of LBAC policies. It relies on functionalities provided by a specialized privacy-aware Location Middleware to collect information about the positions of the User involved in the access control decision process.
- *Location Middleware (LM).* It represents the core component of the architecture. It manages the low-level communications with the Location Provider and enforces both the privacy preferences of the User and the need of location accuracy requested by the Access Control Engine.
- *Location Provider (LP).* It is the component that manages sensing technologies to provide location measurement of the User to the Location Middleware.

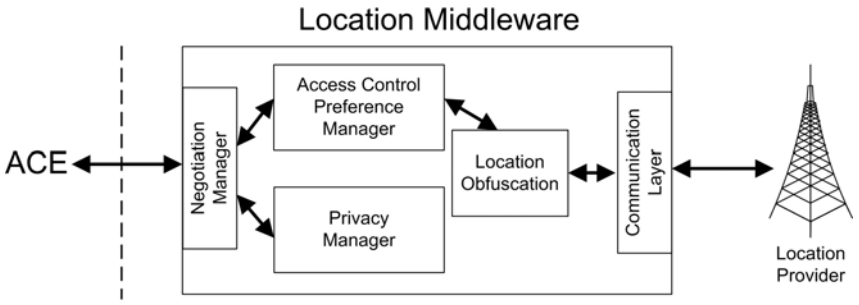

Fig. 3. Location Middleware

The location middleware, whose logical schema is depicted in Fig. 3, includes the following components.

- *Communication Layer.* It manages the communication process with Location Providers by hiding low-level communication details to other components.
- *Negotiation Manager.* It acts as an interface with the Access Control Engine to provide negotiation functionalities regarding service quality and availability based on specific negotiation protocols [34].
- *Access Control Preference Manager.* It manages location service attributes and quality parameters by interacting with the Location Obfuscation component.
- *Location Obfuscation.* It applies obfuscation techniques to location measurements for protecting location privacy of users.
- *Privacy Manager.* It manages privacy preferences expressed by users and supports the privacy-aware location-based predicate evaluation.

A key aspect of such a privacy-aware LBAC architecture is the choice of the component in charge of evaluating LBAC predicates. Although LBAC policy evaluation and enforcement are logically provided by the ACE (i.e., the LBAC system), the LBAC predicates evaluation could take place in two different ways:

- *ACE Evaluation:* the ACE requests to the LM location information relevant to the access decision, without communicating the actual LBAC predicate to be evaluated. The returned response from the LM to the ACE is an obfuscated location measurement with associated a relevance value $\mathcal{R}_{Priv}$ that characterizes its accuracy. Given the relevance $\mathcal{R}_{LBAC}$, the ACE evaluates the LBAC predicate. Since $\mathcal{R}_{LBAC}$ represents the minimum accuracy level that the ACE is willing to accept for a certain service provisioning, $\mathcal{R}_{LBAC} \leq \mathcal{R}_{Eval}$ must hold or the evaluation of the location predicate is rejected.
- *LM Evaluation:* the ACE communicates to the LM the actual LBAC predicate and requests its evaluation based on location information managed by

the LM. The returned response from the LM to the ACE is assumed to be a boolean value with associated a relevance value $\mathcal{R}_{Eval}$ that characterizes the accuracy. $\mathcal{R}_{Eval}$ is derived from $\mathcal{R}_{Priv}$ by considering the obfuscated area generated by the LM and the LBAC predicate. The meaning of $\mathcal{R}_{Eval}$ is the reliability of the predicate evaluation, which depends on the accuracy $\mathcal{R}_{Priv}$ of the obfuscated location information. The LM calculates $\mathcal{R}_{Eval}$ as follows:

$$\mathcal{R}_{Eval} = \frac{Area_{Priv \cap LBAC}}{Area_{Priv}} \cdot \mathcal{R}_{Priv} \tag{7}$$

where the scalar factor $\frac{Area_{Priv \cap LBAC}}{Area_{Priv}}$ depends on the degree of overlapping between the areas resulting by the application of the obfuscation techniques to the location measurement of the user and the area specified by the LBAC predicate (i.e., $Area_{Priv \cap LBAC}$). Again, $\mathcal{R}_{LBAC} \leq \mathcal{R}_{Eval}$ must hold.

Both solutions are viable, although well-suited for different sets of requirements. On the one side, the ACE Evaluation provides a clear separation between business-oriented components (i.e., ACE and Business Application) and location services (i.e., LM and LP). In addition, ACE Evaluation assures that the LM never deals with application-dependent predicates and the ACE never releases information about its access control policies. On the other side, LM Evaluation avoids releasing location information to the ACE. In this setting, location information is always managed by LM that becomes the only trusted component of the architecture with regard to location privacy.

6 Conclusions

Information regarding physical locations of users is rapidly becoming easily available for processing by online and mobile location-based services. Combined with novel application opportunities, however, threats to personal privacy are gaining special prominence, as witnessed by recent security incidents targeting privacy of individuals. This chapter has presented the main techniques aimed at protecting location privacy. The chapter has also described a privacy-aware LBAC architecture that integrates users privacy preferences, obfuscation techniques for location privacy protection, and privacy-enhanced location-based access control.

Acknowledgments

This work was partially supported by the European Union within the PRIME Project in the FP6/IST Programme under contract IST-2002-507591, by the Italian Ministry of Research Fund for Basic Research (FIRB) under project RBNE05FKZ2 and by the Italian MIUR under project MAPS.

References

1. Duckham, M., Kulik, L.: Location privacy and location-aware computing. In: Dynamic & Mobile GIS: Investigating Change in Space and Time. Taylor & Francis (2006) 34–51
2. Lee, J.W.: Location-tracing sparks privacy concerns. Korea Times. http://times.hankooki.com, 16 November 2004. Accessed 22 December 2006
3. Foxs News: Man Accused of Stalking Ex-Girlfriend With GPS. http://www.foxnews.com/story/0,2933,131487,00.html, 04 September 2004. Accessed 22 March 2007
4. Beresford, A.R., Stajano, F.: Location privacy in pervasive computing. IEEE Pervasive Computing **2**(1) (2003) 46–55
5. Bettini, C., Wang, X., Jajodia, S.: Protecting privacy against location-based personal identification. In: Proc. of the 2nd VLDB Workshop on Secure Data Management, LNCS 3674, Springer-Verlag (2005)
6. Gruteser, M., Grunwald, D.: Anonymous usage of location-based services through spatial and temporal cloaking. In: Proc. of the 1st International Conference on Mobile Systems, Applications, and Services. (May 2003)
7. Duckham, M., Kulik, L.: A formal model of obfuscation and negotiation for location privacy. In: Proc. of the 3rd International Conference PERVASIVE 2005, Munich, Germany (May 2005)
8. Ciriani, V., De Capitani di Vimercati, S., Foresti, S., Samarati, P.: K-Anonymity. In: Security in Decentralized Data Management. Springer (2007)
9. Samarati, P.: Protecting respondents' identities in microdata release. IEEE Transactions on Knowledge and Data Engineering **13**(6) (2001) 1010–1027
10. Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., Samarati, P.: Supporting location-based conditions in access control policies. In: Proc. of the ACM Symposium on Information, Computer and Communications Security (ASIACCS'06), Taipei, Taiwan (March 2006)
11. Bonatti, P., Samarati, P.: A unified framework for regulating access and information release on the web. Journal of Computer Security **10**(3) (2002) 241–272
12. Marsit, N., Hameurlain, A., Mammeri, Z., Morvan, F.: Query processing in mobile environments: a survey and open problems. In: Proc. of the 1st International Conference on Distributed Framework for Multimedia Applications (DFMA'05), Besancon, France (February 2005)
13. Jajodia, S., Samarati, P., Sapino, M., Subrahmanian, V.: Flexible support for multiple access control policies. ACM Transactions on Database Systems **26**(2) (June 2001) 214–260
14. OASIS: eXtensible Access Control Markup Language (XACML) Version 1.0. http://www.oasis-open.org/committees/xacml. (2003)
15. van der Horst, T., Sundelin, T., Seamons, K., Knutson, C.: Mobile trust negotiation: Authentication and authorization in dynamic mobile networks. In: Proc. of the 8th IFIP Conference on Communications and Multimedia Security, Lake Windermere, England (September 2004)
16. Gedik, B., Liu, L.: Location privacy in mobile systems: A personalized anonymization model. In: Proc. of the 25th International Conference on Distributed Computing Systems (IEEE ICDCS 2005), Columbus, Ohio (June 2005)
17. Gruteser, M., Bredin, J., Grunwald, D.: Path privacy in location-aware computing. In: Proc. of the Second International Conference on Mobile Systems,

Application and Services (MobiSys2004), Boston, Massachussetts, USA (June 2004)
18. Gruteser, M., Liu, X.: Protecting privacy in continuous location-tracking applications. IEEE Security & Privacy Magazine **2**(2) (March-April 2004) 28–34
19. Ho, B., Gruteser, M.: Protecting location privacy through path confusion. In: Proc. of IEEE/CreateNet International Conference on Security and Privacy for Emerging Areas in Communication Networks (SecureComm), Athens, Greece (September 2005)
20. Mokbel, M., Chow, C.Y., Aref, W.: The new casper: Query processing for location services without compromising privacy. In: Proceedings of the 32nd International Conference on Very Large Data Bases, Korea (September 2006) 763–774
21. Beresford, A.R., Stajano, F.: Mix zones: User privacy in location-aware services. In: Proc. of the 2nd IEEE Annual Conference on Pervasive Computing and Communications Workshops (PERCOMW04). (2004)
22. Hauser, C., Kabatnik, M.: Towards Privacy Support in a Global Location Service. In: Proc. of the IFIP Workshop on IP and ATM Traffic Management (WATM/EUNICE 2001), Paris, France (March 2001)
23. Geopriv: Geographic Location/Privacy. http://www.ietf.org/html.charters/geopriv-charter.html. (September 2006)
24. Hong, D., Yuan, M., Shen, V.Y.: Dynamic privacy management: a plug-in service for the middleware in pervasive computing. In: Proc. of the 7th International Conference on Human Computer Interaction with Mobile Devices & Services (MobileHCI'05), Salzburg, Austria (2005)
25. Langheinrich, M.: A privacy awareness system for ubiquitous computing environments. In Borriello, G., Holmquist, L.E., eds.: Proc. of the 4th International Conference on Ubiquitous Computing (Ubicomp 2002). (September 2002) 237–245
26. W3C: Platform for privacy preferences (p3p) project. http://www.w3.org/TR/P3P/. (April 2002)
27. Duckham, M., Kulik, L.: Simulation of obfuscation and negotiation for location privacy. In: Proc. of Conference On Spatial Information Theory (COSIT 2005). (September 2005) 31–48
28. Openwave: Openwave Location Manager. http://www.openwave.com/. (2006)
29. Bellavista, P., Corradi, A., Giannelli, C.: Efficiently managing location information with privacy requirements in wi-fi networks: a middleware approach. In: Proc. of the International Symposium on Wireless Communication Systems (ISWCS'05), Siena, Italy (September 2005)
30. Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., Samarati, P.: Managing privacy in LBAC systems. In: Proc. of the Second IEEE International Symposium on Pervasive Computing and Ad Hoc Communications (PCAC-07), Niagara Falls, Canada (May 2007)
31. Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., Samarati, P.: A middleware architecture for integrating privacy preferences and location accuracy. In: Proc. of the 22nd IFIP TC-11 International Information Security Conference (SEC 2007), Sandton, South Africa (May 2007)
32. Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., Samarati, S.: Location privacy protection through obfuscation-based techniques. In: Proc. of the 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security, Redondo Beach, CA, USA (July 2007)

33. Atluri, V., Shin, H.: Efficient enforcement of security policies based on tracking of mobile users. In: Proc. of the 20th Annual IFIP WG 11.3 Working Conference on Data and Applications Security, Sophia Antipolis, France (July-August 2006) 237–251
34. Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., Samarati, P.: Location-based metadata and negotiation protocols for LBAC in a one-to-many scenario. In: Proc. of the Workshop on Security and Privacy in Mobile and Wireless Networking (SecPri_MobiWi 2006), Coimbra, Portugal (May 2006)