

A Middleware Architecture for Integrating Privacy Preferences and Location Accuracy

Claudio Agostino Ardagna, Marco Cremonini, Ernesto Damiani,
Sabrina De Capitani di Vimercati, and Pierangela Samarati

Università degli Studi di Milano - Dipartimento di Tecnologie dell'Informazione
Via Bramante 65, 26013 Crema (CR) - Italy
{ardagna, cremonini, damiani, decapita, samarati}@dti.unimi.it

Abstract. Location-Based Access Control (LBAC) systems support the evaluation of conditions on locations in the enforcement of access control policies. The ability to evaluate conditions on a set of authorized locations has a number of well-known advantages, including enriching access control expressiveness. However, when locations are used in combination with personal identities, users privacy must be considered. In this paper, we describe a solution to integrate a LBAC system with privacy-enhanced techniques based on location obfuscation. Our solution is based on a privacy-aware middleware component that explicitly addresses the trade-off between users privacy and location accuracy by satisfying preferences set by users and maximizing the quality of location information released to LBAC systems.

1 Introduction

In ubiquitous and mobile computing, user position is a fundamental attribute for managing location-based applications. Access to location information is achieved through a variety of sensor technologies, which recently enjoyed a relevant boost in term of precision and reliability. As a secondary effect of improved location capabilities, protection of user location privacy has become one of the hottest and most critical topics. In this paper, we address location privacy in the framework of location-based services (LBSs). Specifically, we consider *Location-Based Access Control* (LBAC) systems, which support access control policies based on the physical locations of users. Within the class of applications based on LBAC, some necessarily require the best location accuracy for their provision, like those working in mission-critical environments or aimed at providing emergency services. In these cases, privacy concerns are of lesser importance and must be treated specifically for each particular situation. Differently, many other applications based on LBAC could accept location information with sub-optimal accuracy and still offer an acceptable quality of service. In these cases, one of the most critical LBAC issue is to find a balance between *location accuracy* and *location privacy*, dealing with requirements from both business applications and user privacy. The expressive power and the granularity of LBAC

Please use the following format when citing this chapter:

Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., and Samarati, P., 2007, in IFIP International Federation for Information Processing, Volume 232. New Approaches for Security, Privacy and Trust in Complex Environments, eds. Venter, H., Eloff, M., Labuschagne, L., Eloff, J., von Solms, R., (Boston: Springer), pp. 313–324.

policies, in fact, heavily depend on the accuracy of the locations of users, and the disclosure of fine-grained location information to the LBAC enforcement engine must comply with user's privacy preferences and regulations.

Generally speaking, location-based services require two separate contractual agreements: *i*) between the user and a telecommunication company¹ acting as (or on behalf of) location provider, and *ii*) between the location provider and the application requiring LBAC policies. This dual agreement is critical because, as a generic subscriber to the mobile phone network, an individual may want her privacy strictly preserved, while, as a user of location-based services she may want the service provider to handle very accurate location information to receive best-quality service. To address this issue, we introduce location obfuscation techniques to protect the privacy of the location of users and a distributed architecture (built around a privacy-aware location middleware) decoupling business applications and LBAC policy enforcement from location providers. This way, location middleware can effectively and securely manage a trade-off between accuracy and privacy. The remainder of this paper is organized as follows. Section 2 describes related work. Section 3 illustrates the basic features of our LBAC system and the obfuscation techniques used to achieve user privacy. Section 4 illustrates our privacy-aware architecture, discusses different solutions for evaluating location-based predicates, and shows the working of our privacy-aware middleware. Finally, Section 5 gives our conclusions.

2 Related work

The definition of LBAC systems is an emerging research area that has not been fully investigated yet. Some papers recently present architectures, designed for pervasive environments, that incorporate mobile data for security management [2]. Others consider location information as a resource to protect against unauthorized access [3, 4, 6]. Beresford and Stajano [3] refine a method, called *mix zones*, to enhance privacy in location-based services. Their proposal uses a trusted middleware that anonymizes location information. Bettini et al. [4] present an investigation of the privacy issues raised by a location-based services scenario. Duckham and Kulik [6] investigate obfuscation techniques for protecting the privacy of the locations of users.

Other works propose special-purpose *location middleware* for managing interactions between applications and location providers, while maximizing the quality of service (QoS) [9, 10, 11]. Typically, in these proposals the location middleware *i*) receives requests from LBS components asking for location information, *ii*) collects users locations from a pool of location providers, and *iii*) produces an answer. Naguib et al. [9] present a middleware framework, called *QoSDREAM*, for managing context-aware multimedia applications. Nahrstedt

¹ This is true regardless of the specific location technology used. For instance, satellite location information like GPS is made available to applications via the mobile network.

et al. [10] present a QoS middleware for ubiquitous computing environments aimed at maximizing the QoS of distributed applications. Ranganathan et al. [11] present a middleware that provides a clear separation between business applications and location detection technologies. They also address the issue of managing location data from heterogeneous location technologies.

Although several middleware components supporting communication and negotiation between location services and applications have been presented, only a few proposals try to integrate service quality and privacy protection. For instance, Myles et al. [8] propose an architecture aimed at preserving privacy in location-based services. The architecture is based on a middleware managing the interactions between location-based applications and location providers and on the definition of policies for data release. Hong et al. [7] present an extension of the P3P language for representing user privacy preferences for context-aware applications. The main drawback of their solution is that users are seldom willing to directly manage complex preference policies. By contrast, in our approach users have only to specify a few simple and intuitive parameters. Similarly to [7, 8], our work defines an architecture centered on a middleware component aimed at balancing service accuracy and privacy protection requirements. Differently from other works, our work focuses on some obfuscation-based techniques that degrade a location accuracy and introduces a formal location privacy estimator, called *relevance*.

3 Privacy and LBAC systems

The definition of location-based conditions and their management is the first step towards the development of a privacy-aware LBAC architecture. We identify three main classes of conditions to be included in access control policies and whose evaluation is actually possible with current technology: *movement-based*, *position-based*, and *interaction-based* [2]. Starting from these classes, a set of predicates corresponding to specific conditions can be defined. For instance, predicate $\text{inarea}(\text{user_term}, \text{area_term})$ is a binary position predicate, where the first parameter represents a user and the second parameter is a geographical area. The predicate semantics is to evaluate whether a user is located within a specific area (e.g., a city, a street, a building). When evaluating location-based predicates, however, we need to consider that location-based information is radically different from other context-related knowledge inasmuch it is both *approximate* (all location systems have a margin of error) and *time-variant* (location is subject to fast changes, especially when the user is in motion).

To accommodate these peculiar characteristics of location-based predicates, we introduce the notion of *relevance* as the estimator of the accuracy of all location-based measurements and evaluations. A relevance is a number $\mathcal{R} \in [0, 1]$ that assumes value 0 when there is no accuracy in the location-based evaluation/measurement, value 1 for full accuracy, and values in $(0, 1)$

to represent various degrees of accuracy. Accordingly, the guaranteed location privacy is $(1-R)$. A LBAC system has to manage the following relevance values.

LBAC relevance (\mathcal{R}_{LBAC}). The *minimum* accuracy required by business applications for a user location measurement or for a location-based predicate evaluation. It represents the lowest acceptable quality of a location service.

Privacy relevance (\mathcal{R}_{Priv}). The *maximum* location relevance accepted by a user for her location information. It represents the highest acceptable location accuracy according to user's privacy preferences.

Technological relevance (\mathcal{R}_{Tech}). The measurement accuracy provided by a location provider given a certain mobile technology and environment.

All these relevance values represent the degree of accuracy related to a location measurement. \mathcal{R}_{Tech} and \mathcal{R}_{LBAC} are assumed to be given, while \mathcal{R}_{Priv} is the result of the application of suitable *obfuscation techniques*. Our goal is to apply an obfuscation technique to location measurements in such a way that the following relation holds: $\mathcal{R}_{LBAC} \leq \mathcal{R}_{Priv} \leq \mathcal{R}_{Tech}$. Given a location measurement with relevance \mathcal{R}_{Tech} , some transformations are applied to make it less accurate, so that privacy requirements can be met. The resulting location measurement retains a level of relevance (\mathcal{R}_{Priv}), which has to be greater than \mathcal{R}_{LBAC} to be meaningful for LBAC enforcement.

3.1 Location obfuscation and user privacy

Obfuscation techniques applied to a location measurement increase the uncertainty of a user location by degrading its accuracy. In this work, we shall consider a planar (2-D) coordinate space for locations. Also, since the result of each location measurement is necessarily affected by an error, a *spatial area* is always returned, rather than a single point. We introduce two working assumptions: *i*) the area returned by a location measurement is circular, which is the actual shape resulting from many location technologies [5]; *ii*) the distribution of measurement errors within a returned area is uniform. This last assumption increases accuracy and precision, which are the main requirements for LBAC predicate evaluation. According to these two assumptions, we formally define a location measurement and the associated error as follows.

Definition 3.1 (Location measurement) *A location measurement of a user u is a circular area, denoted $Area(r, x_c, y_c)$, centered on the geographical coordinates (x_c, y_c) and with radius r , which includes the real position of u .*

Definition 3.2 (Uniform distribution) *Given a location measurement $Area(r, x_c, y_c)$, the distribution is uniform if and only if the corresponding probability density function (pdf) $f_r(x, y)$ is:*

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

Since our main goal in obfuscating a location measurement is to select an area corresponding to a given relevance \mathcal{R}_{Priv} , we need to better specify how \mathcal{R}_{Priv} is calculated. In a real world scenario, it is very unlikely that a user could explicitly specify such a value (what would a 0.6 relevance exactly mean?). Many proposals in the location privacy field assume that users specify their privacy preferences in terms of intuitive parameters such as *minimum distance* [6]. For instance, a user can require the radius of the location area to be at least 100 meters. In this case, obfuscation is achieved by increasing measurement granularity. Although minimum distance is easy to understand and implement, it has a severe drawback: an absolute distance value is only meaningful when related to a specific application context. In the previous example, the value of 100 meters is well suited to applications that provide touristic information to a user walking in a city center. Location-based applications working, for example, in smaller contexts, as inside a production plant, are likely to become ineffective if the granularity is 100 meters. Also, 100 meters can be insufficient for preserving user privacy in high sensitive contexts.

A different (and equally intuitive) way for users to specify privacy requirements is for a relative degradation of the measure with respect to the location accuracy (i.e., \mathcal{R}_{Tech}). In our approach, privacy preferences are therefore defined through a simple index $\lambda \in [0, \infty]$ that represents the *privacy rate* in terms of degradation applied to location accuracy. For instance, if a user asks no privacy, then $\lambda = 0$. If a user asks total privacy, $\lambda \rightarrow \infty$. Normally, a user may ask that the accuracy of her location must be decreased by a certain rate, like 100%, which implies $\lambda = 1$, or 200%, which implies $\lambda = 2$.

Both minimum distance d and rate λ are easy to specify for users. Among the two, λ is the more general index because independent from the specific application context and measurement unit.²

3.2 Obfuscation by scaling the radius

The first and most obvious technique for obfuscating a location measurement is to scale the radius of the circular area. The obfuscation effect directly derives from Definition 3.2: $\forall r, r_u : r < r_u : f_r(x, y) > f_{r_u}(x, y)$. Fig. 1(a) shows the effect of obfuscation by scaling the radius, where the circular area of radius r is the area returned by a sensing technology and the area of radius r_u is the obfuscated area. The relevance \mathcal{R}_{Priv} of the obfuscated location is calculated by dividing the pdf of the obfuscated area by the pdf of the original area multiplied by \mathcal{R}_{Tech} :

$$\text{given } r, r_u : r < r_u, \mathcal{R}_{Priv} = \frac{r^2}{r_u^2} \mathcal{R}_{Tech} \quad (1)$$

² Parameter λ depends on the accuracy of each measurement realized with a specific location technology. In this paper, we assume that a single location technology is used and users are aware of the best accuracy that the technology can achieve. We plan to develop a more general approach in future work.




Fig. 1. Obfuscation by scaling the radius (a) and by shifting the center (b)

Otherwise, if the rate λ is used to specify the privacy preference, the new radius r_u can be derived as follows:

$$\text{given } \lambda \geq 0: \mathcal{R}_{Priv} = (\lambda + 1)^{-1} \mathcal{R}_{Tech}, \quad r_u = r\sqrt{\lambda + 1}. \quad (2)$$

3.3 Obfuscation by center-shifting

Shifting the center of the location area is another way of decreasing its accuracy. The obfuscated area is derived from the original area either by setting the distance d between the two centers to the value specified by the user or by deducing d from rate λ . Let $Area(r, x_c + \Delta x, y_c + \Delta y)$ be the obfuscated area and suppose that the distance d is greater than or equal to $2r$. In this situation, the probability that the obfuscated area contains the real position of the user (i.e., (x_u, y_u)) is zero, that is, $P((x_u, y_u) \in Area(r, x_c + \Delta x, y_c + \Delta y)) = 0$. Otherwise (i.e., $0 < d < 2r$), $0 < P((x_u, y_u) \in Area(r, x_c + \Delta x, y_c + \Delta y)) < 1$.

The privacy gain can be quantitatively measured by considering the intersection of the original and the obfuscated area, 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 between the two centers. In particular, if $d = 0$, there is no privacy gain; if $d \geq 2r$, there is full privacy; and if $0 < d < 2r$, there is an increment of privacy.

To derive the actual obfuscated area, the angle θ illustrated in Fig. 1(b) must be chosen too. With regard to θ , however, there is no meaning for a user to specify it, so it must be defined by the component in charge of obfuscating the location measurement. The first and obvious choice is to randomly choose θ , because in general all its values are equivalent with respect to user privacy preferences. However, in the next section, we will discuss how a reasonable choice of this parameter can be made to maximize the relevance associated with location-based evaluations, still preserving user preferences.

\mathcal{R}_{Priv} can be derived from the ratio of the intersection $Area_{Tech \cap Priv}$ over the obfuscated area as follows.

$$\mathcal{R}_{Priv} = (\lambda + 1)^{-1} \cdot \mathcal{R}_{Tech} = \frac{Area_{Tech \cap Priv}}{Area(r, x_c + \Delta x, y_c + \Delta y)} \cdot \mathcal{R}_{Tech} \quad (3)$$

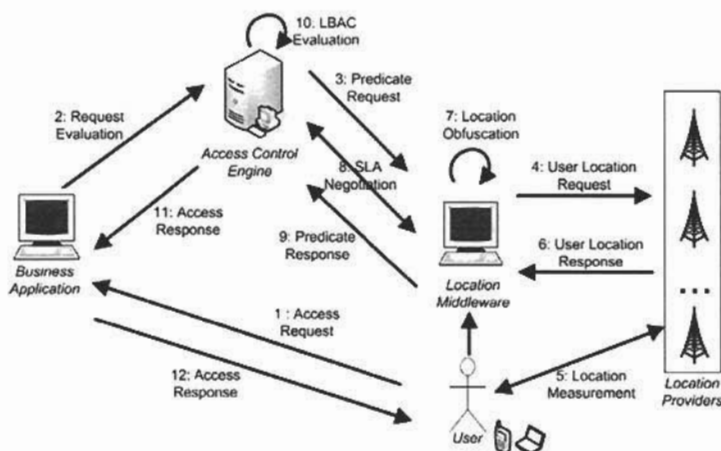


Fig. 2. Privacy-Aware LBAC Architecture

4 A privacy-aware LBAC architecture

The above concepts and techniques are at the base of the definition of our privacy-aware LBAC architecture. The logical components of the architecture are showed in Fig. 2 and can be summarized as follows.

User. Individual to be located through her mobile terminal.

Business application. Customer-oriented application that provides resources protected by LBAC policies.

Access Control Engine (ACE). A component that stores and enforces LBAC policies. For the enforcement, it requests location services and information from the Location Middleware.

Location Providers (LPs). Components using location sensing technologies to provide location measurements.

Location Middleware (LM). The entity that interacts with different LPs and provides location services to the ACE. It has to satisfy users privacy preferences and ACEs location accuracy needs.

Communications among these components are performed via request/response message exchanges. Basically, the interaction flow can be logically partitioned in six macro-operations: *i) initialization*, when user preferences and LBAC policies are defined; *ii) location information retrieval*, when LM collects user location information through a communication process with multiples LPs; *iii) SLA negotiation*, when a Service Level Agreement (SLA) specifying QoS attributes and corresponding service cost is agreed between an ACE and a LM; *iv) location obfuscation*, when obfuscation techniques are used to comply with both user preference and LBAC accuracy; *v) LBAC evaluation*, when the LBAC policies are evaluated; and finally *vi) access decision*, when the access request is granted or denied.

4.1 LBAC predicates evaluation

A major design issue for our privacy-aware LBAC architecture is related to the component in charge of evaluating LBAC predicates. Two choices are possible, which deeply affect how privacy is guaranteed.

ACE evaluation: The ACE asks users locations to LM without disclosing LBAC predicates.

LM evaluation: The ACE sends to LM a LBAC predicate for evaluation and receives a boolean answer and a relevance value.

Both choices are viable and well-suited for different set of requirements. On one side, *ACE evaluation* enforces a clear separation between applications and location services because the location service infrastructure (i.e., LMs and LPs) never deals with application-dependent location-based predicates. On the other side, *LM evaluation* avoids the exchange of user locations, although obfuscated, with applications. This second choice is also more flexible in business terms. For instance, an ACE can subscribe to a location service for a specific set of location predicates, and select different QoS according to different needs (e.g., different accuracy levels). The LM could then differentiate prices according to service quality.

Since the analysis presented so far has implicitly assumed the ACE evaluation scenario (i.e., the ACE component receives an obfuscated area with a given \mathcal{R}_{Priv} value), we now describe how LM evaluation is carried out. The main difference is that now LM returns an answer for the LBAC predicate evaluation together with a relevance of that answer, which we call \mathcal{R}_{Eval} . This relevance is derived from \mathcal{R}_{Priv} by considering both the obfuscated area and the area specified into the LBAC predicate. Since the ACE component requires a minimum acceptable relevance \mathcal{R}_{LBAC} , $\mathcal{R}_{Eval} \geq \mathcal{R}_{LBAC}$ must hold. For instance, let $\text{inarea}(\text{JohnID}, \text{Room1})$ be the predicate that the ACE component sends to the LM component, which asks whether the user *JohnID* is in room *Room1*. LM calculates \mathcal{R}_{Eval} as follows:

$$\mathcal{R}_{Eval} = \frac{\text{Area}_{Priv \cap LBAC}}{\text{Area}_{Priv}} \cdot \mathcal{R}_{Priv} \quad (4)$$

where the scalar factor depends on the intersection, denoted $\text{Area}_{Priv \cap LBAC}$, between the obfuscated area and the area specified by the LBAC predicate.

There is, however, a subtlety to consider when center-shifting obfuscation is applied. As noted in Section 3.1, there are infinite values of angle θ that could be chosen, all equivalent with respect to the \mathcal{R}_{Priv} value. When the LBAC predicate is evaluated, however, the choice of θ is relevant, because according to the position of the obfuscated area, the \mathcal{R}_{Eval} value may change. This requires the following additional constraint:

$$\mathcal{R}_{Eval} \leq \frac{\text{Area}_{Tech \cap LBAC}}{\text{Area}_{Tech}} \cdot \mathcal{R}_{Tech} \quad (5)$$

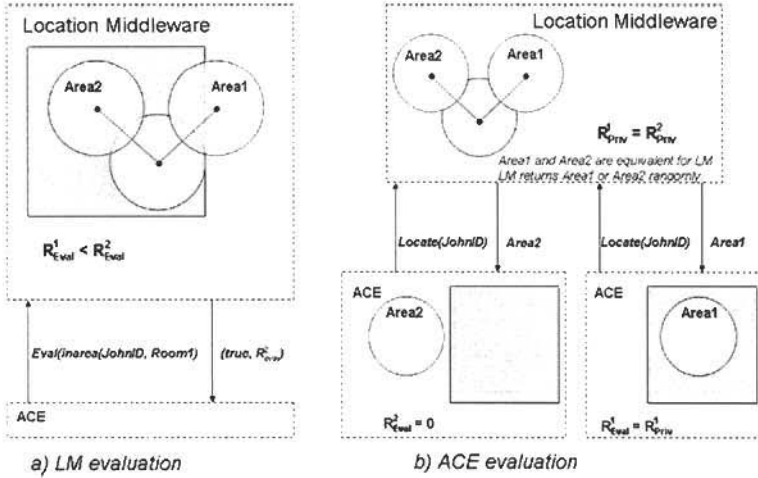


Fig. 3. An example of LM evaluation (a), and ACE evaluation (b)

The rationale for this constraint is to avoid the case of a relevance \mathcal{R}_{Eval} derived from \mathcal{R}_{Priv} that is greater than the one that would have provided the original area with relevance \mathcal{R}_{Tech} . In other terms, areas must not be manipulated with obfuscation techniques just to increase the odds of satisfying LBAC quality requirements. This case would be made possible by shifting the center in such a way that, for example, the obfuscated area is completely included into the area specified by the predicate (*Room1*, in our example), while the original area is just partially included. Our constraint ensures that, given an infinite set Θ of angles, a set $\Theta_f \subseteq \Theta$ is generated, containing all angles $\theta_1 \dots \theta_n$ that produce a relevance \mathcal{R}_{Eval} at most equals to the relevance produced by considering the original area.

When center-shifting obfuscation is adopted, the ACE vs LM choice has a significant impact. To illustrate, consider the examples in Fig. 3(a) and Fig. 3(b) that show the evaluation of predicate $inarea(JohnID, Room1)$ in case of LM evaluation and of ACE evaluation, respectively. Here, *Area1* and *Area2* are two possible obfuscated areas.

If LM evaluation is performed, LM computes \mathcal{R}_{Eval} from (4) and is able to establish an ordering among obfuscated areas according to the different values of \mathcal{R}_{Eval} . In our example, it is easy to see that relevance \mathcal{R}_{Eval}^2 resulting from *Area2* is greater than relevance \mathcal{R}_{Eval}^1 resulting from *Area1*. This information is important for the provision of the location service, because when returned to ACE, the value \mathcal{R}_{Eval} is matched with \mathcal{R}_{LBAC} , the minimum relevance required for LBAC evaluation. The best strategy for LM is therefore to select the angle θ that produces the obfuscated area that, given \mathcal{R}_{Priv} , maximizes \mathcal{R}_{Eval} .

If ACE evaluation is in place, LM does not calculate any \mathcal{R}_{Eval} (i.e., ACE does not communicate the location predicate under evaluation), and it can only




Fig. 4. Location Middleware

select randomly one value for θ among all those that produce an obfuscated area with the same \mathcal{R}_{Priv} . In this way, random selection of the obfuscated area (in our example, *Area1* or *Area2*) may cause an unpredictable result during ACE evaluation, ranging from relevance equal to zero (e.g., when *Area2* in Fig. 3(b) is returned) to relevance equal to \mathcal{R}_{Priv} (e.g., when *Area1* in Fig. 3(b) is returned). As a consequence, also the matching with the condition over \mathcal{R}_{LBAC} results in random rejection or acceptance of the predicate evaluation. Therefore, center-shifting obfuscation is incompatible with ACE evaluation. This result supports architectures including location middleware capable of autonomously evaluating LBAC predicates.

4.2 The privacy-aware middleware

As mentioned in Section 2, currently available middleware components are mostly in charge of managing interactions between applications and location providers, and managing communication and negotiation protocols aimed at maximizing the QoS. Instead, in our approach the privacy-aware middleware has to find a balance between users privacy and location-based services accuracy. To this end, our LM is responsible both for the obfuscation of user locations and for the location-based predicates evaluation. As shown in Fig. 4, LM is functionally divided into the following five logical components.

Communication Layer. It manages the communication process with LPs. Hides low-level communication details to other components.

Negotiation Manager. It acts as an interface with ACE. It provides negotiation functionalities and implements the negotiation protocols [1].

Access Control Preference Manager. It manages location service attributes and quality by interacting with the Location Obfuscation component.

Location Obfuscation. It applies obfuscation techniques for users privacy.

Privacy Manager. It manages privacy preferences and location-based predicate evaluation.

As an example of the LM operations, assume that user John subscribes to the LM by setting his privacy preference to $\lambda = 0.2$, which is meant to degrade location accuracy by 20%. After that, John uses a business application

that adopts LBAC policies. In particular, one of these LBAC policies states that when a user is in *Room1*, she gains the access to an online financial service. Additionally, the ACE component is set to require a minimum evaluation relevance $\mathcal{R}_{LBAC} = 0.7$. To grant or deny John's access to the online financial service, the ACE sends to LM a predicate evaluation request for predicate $\text{inarea}(\text{JohnId}, \text{Room1})$ together with relevance \mathcal{R}_{LBAC} . The LM asks to the LP (for simplicity, suppose that only one LP is available) the John's position and receives as an answer a circular area together with a technology relevance $\mathcal{R}_{Tech} = 0.9$, representing the accuracy of the measurement. At this point, LM must obfuscate the location, for example, by shifting the center. It calculates $\mathcal{R}_{Priv} = (\lambda + 1)^{-1} \mathcal{R}_{Tech} = 0.75$. Among all possible values of angle θ that produce an obfuscated area with $\mathcal{R}_{Priv} = 0.75$, LM has to select the obfuscated area that maximizes the corresponding relevance \mathcal{R}_{Eval} computed as in (4) and that satisfies the restriction defined in (5). For simplicity, we only consider *Area1* and *Area2* illustrated in Fig. 3. *Area2* falls completely into the square greyed box representing the geometry and position of *Room1*, so $\mathcal{R}_{Eval}^2 = \mathcal{R}_{Priv} = 0.75$. *Area1*, instead, is partially overlapped with the grey box, so $\mathcal{R}_{Eval}^1 < \mathcal{R}_{Priv}$. Both satisfy the restriction defined in (5), therefore LM can return to ACE a *true* evaluation of the inarea predicate together with $\mathcal{R}_{Eval}^2 = 0.75$. Finally, the ACE can proceed in the enforcement of the LBAC policy having its location predicate positively evaluated, that is, the corresponding boolean value is *true* and the evaluation relevance is greater than $\mathcal{R}_{LBAC} = 0.7$.

It is important to highlight that the architecture of our location middleware can be extended to include the important case of users setting *multiple privacy preferences* according to different contexts. For instance, there could be users wishing to set: no privacy preferences for location services dedicated to the social network of their relatives and close friends; a certain level of privacy for business location services aimed at helping to find point of interests (e.g., shops, or monuments), and for location services whose goal is to find their position while at work; and strong privacy requirements in high sensitive contexts.

5 Conclusions

In this paper, we presented an architecture built around a location middleware for evaluating LBAC predicates. We have showed a solution that supports the critical issue of striking a balance between accuracy and privacy requirements. To the best of our knowledge, this is the first middleware solution that smoothly manages such aspects of a LBAC infrastructure through different obfuscation techniques and an uniform index representing a common estimator for both quality and privacy requirements. Future work to be carried out includes extending our architecture to fully support the multiple privacy preferences scenario and enriching LM with the ability to deal with contextual information.

Acknowledgments

This work was partially supported by the European Union within the PRIME Project 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. C.A. Ardagna, M. Cremonini, E. Damiani, S. De Capitani di Vimercati, and P. Samarati. 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*, Coimbra, Portugal, May 2006.
2. C.A. Ardagna, M. Cremonini, E. Damiani, S. De Capitani di Vimercati, and P. Samarati. 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.
3. A. R. Beresford and F. Stajano. Mix zones: User privacy in location-aware services. In *Proc. of the 2nd IEEE Annual Conference on Pervasive Computing and Communications Workshops (PERCOMW'04)*, Orlando, Florida, March 2004.
4. C. Bettini, X.S. Wang, and S. Jajodia. Protecting privacy against location-based personal identification. In *Proc. of the 2nd VLDB Workshop on Secure Data Management*, Trondheim, Norway, September 2005.
5. E. Damiani, M. Anisetti, and V. Bellandi. Toward exploiting location-based and video information in negotiated access control policies. In *Proc. of the 1st International Conference on Information Systems Security (ICISS 2005)*, Kolkata, India, December 2005.
6. M. Duckham and L. Kulik. A formal model of obfuscation and negotiation for location privacy. In *Proc. of the 3rd International Conference on Pervasive Computing*, Munich, Germany, May 2005.
7. D. Hong, M. Yuan, and V. Y. Shen. 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*, Salzburg, Austria, September 2005.
8. G. Myles, A. Friday, and N. Davies. Preserving privacy in environments with location-based applications. *IEEE Pervasive Computing*, 2(1):56–64, 2003.
9. H. Naguib, G. Coulouris, and S. Mitchell. Middleware support for context-aware multimedia applications. In *Proc. of the IFIP TC6 / WG6.1 3rd International Working Conference on New Developments in Distributed Applications and Interoperable Systems*, Deventer, The Netherlands, September 2001.
10. K. Nahrstedt, D. Xu, D. Wichadakul, and B. Li. QoS-aware middleware for ubiquitous and heterogeneous environments. *IEEE Communications Magazine*, pages 140–148, November 2001.
11. A. Ranganathan, J. Al-Muhtadi, S. Chetan, R. H. Campbell, and M. D. Mickunas. Middlewhere: A middleware for location awareness in ubiquitous computing applications. In *Proc. of the ACM/IFIP/USENIX 5th International Middleware Conference (Middleware 2004)*, Toronto, Ontario, Canada, October 2004.