A Privacy Risk Assessment Methodology for Location-Based Systems

Jesús Friginal, Jérémie Guiochet and Marc-Olivier Killijian

LAAS-CNRS, 7 Avenue du Colonel Roche
31400 Toulouse Cedex, France
{jesus.friginal,jeremie.guiochet,marco.killijian}@laas.fr

Mobiquitous systems are gaining more and more weight in our daily lives. They are becoming a reality from our home and work to our leisure. The use of Location-Based Services (LBS) in these systems is increasingly demanded by users. Yet, while on one hand they enable people to be more "connected", on the other hand, they may expose people to serious privacy issues. The design and deployment of Privacy-Enhancing Technologies (PETs) for LBS has been widely addressed in the last years. However, strikingly, there is still a lack of methodologies to assess the risk that using LBS may have on users' privacy (even when PETs are considered). This paper presents a privacy risk assessment methodology to (i) identify (ii) analyse, and (iii) evaluate the potential privacy issues affecting mobiquitous systems. The feasibility of each step of our methodology is illustrated through an innovative case study of dynamic carpooling.

Key words: privacy, risk assessment, location-based systems

1 Introduction

After the successful development of positioning technology, such as GPS, and the rise of infrastructureless wireless networks, such as ad hoc networks, mobile and ubiquitous (mobiquitous) systems have become the spearhead sector of the communication industry. The vast deployment of myriads of sensors and the rapid growth in the number of mobile devices per person is providing enormous opportunities to create a new generation of innovative Location-Based Services (LBS) addressed to improve the welfare of our society. LBS are used in a variety of contexts, such as health, entertainment or work, like discovering the nearest cash machine, parking parcel, or getting personalised weather services.

Parallel to this revolution, numerous studies reveal potential privacy breaches in the use of these services given the sensitivity of the collected information, and how it is stored and exchanged [4]. From a privacy viewpoint, the main characteristic of LBS systems is that, apart from personal identity, users' location becomes a new essential asset to protect. A recent study from MIT [1] showed that 4 spatio-temporal points (approximate places and times), were enough to unequivocally identify 95% people in a mobility database of 1.5M users. The study shows that these constraints hold even when the resolution of the dataset is low. Therefore, even coarse or blurred information provides little anonymity. However, very few users are aware of the implications that a misuse of their

location information may have on their privacy, and the potential consequences on their security and safety [2]. Tackling this issue becomes critical given the increasingly variety of attacks that may impact the privacy of users [4].

By the time being, there is a range from simplistic on/off switches to sophisticated Privacy-Enhancing Technologies (PETs) using anonymisation techniques [10]. Today, few LBS offer such PETs, e.g., Google Latitude offers an on/off switch that allows to stick one's position to a freely definable location. Another set of techniques include location obfuscation, which slightly alter the location of the users in order to hide their real location while still being able to represent their position and receive services from their LBS provider. However, such efforts remain questionable in practice while suitable techniques to guarantee acceptable levels of risk remain unavailable. Consequently, the confident use of LBS requires not only the development of PETs, but also the definition of methodologies and techniques to assess and treat the privacy risk associated to LBS solutions. There exist many and various challenges in the deployment of LBS, but the need for identifying the risk related to the processing of personal data before determining the appropriate means to reduce them, is without doubt, one of the most important in the domain of LBS. Unfortunately, to date there is an absence of methodologies to adequately approach this problem. The risk assessment proposals found in standards [7, 11] are so generic, that are really difficult to map to privacy and, even more to the domain of LBS.

Therefore, the main goal of this paper is to design a methodology to assess the risk related to the lack of privacy of LBS. The rest of this paper is structured as follows. Section 2 shows the lack of techniques and guidelines to assess the privacy risk on LBS. Section 3 introduces our privacy risk assess methodology for LBS. Section 4 presents a case study based on dynamic carpooling to show the usefulness of our methodology. Finally, Section 5 closes the paper.

2 Privacy in risk standards

The concept of risk was first introduced in safety critical systems, but is now widely used in many domains, including information technology. Indeed, users, environment and organizations could be faced to the harm induced by the use of a new technology. The generic standard ISO/IEC-Guide73 [8] defines the risk as the combination of the probability of an event and its consequence. This definition had to be adapted in the domain of security, where risk is defined as the "potential that a given threat will exploit vulnerabilities of an asset or group of assets and thereby cause harm to the organization" [7]. In this definition, the classic notion of *probability* of an event has been replaced by "potentiality" given the difficulty to estimate such a probability. The concept of *consequence* was also refined into "harm to the organization". The identification, analysis and evaluation of the risk, is defined in many standards and international directives as *risk assessment* within the risk management process, as Figure 1 shows.

ISO standards such as [7] or regulatory documents produced by NIST [11], deal with security risk assessment and treatment. Unfortunately, there is no privacy ISO/IEC standard dedicated to privacy risk assessment. In order to protect

Fig. 1. Risk management process adapted from ISO 27005 [7].

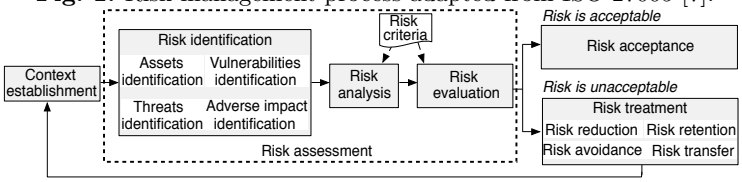

user data, the European Commission unveiled in 2012 a draft [3] that will supersede the Data Protection Directive 95/46/EC. It is mainly about openness and obligations of organizations managing users personal data, and it does not include methods and techniques to analyse and assess the risks. A similar approach is presented in the USA Location Privacy Protection Act of 2012 (S.1233), in order to regulate the transmission and sharing of user location data. As in the European Directives, this bill specifies the collecting entities, the collectable data and its usage, but no analysis techniques or methods are proposed, and much less in the domain of LBS.

3 A privacy risk assessment methodology for LBS

The privacy risk assessment methodology introduced in this Section systematises the identification, analysis and evaluation of privacy risks for LBS.

3.1 Process overview and main concepts

Basically, our methodology addresses the privacy risk assessment following the framework of Figure 1. The first step, risk identification involves the identification of risk sources, events, their causes and their potential consequences. The level of risk is then estimated in the risk analysis step. For this step, we will introduce metrics as risk criteria for the privacy risk analysis. To clearly illustrate the concepts used in our methodology, Figure 2 shows the general meta-model concerning the entities and relations involved in the process. It is worth noting that our meta-model should not be seen as a closed proposal. Instead, it is a flexible abstraction interpreting privacy risk in a generic way for LBS.

As seen in Section 2, risk is generally defined as the combination of two factors: the potential occurrence of threats, and the consequences that such

Fig. 2. General meta-model for privacy risk assessment.

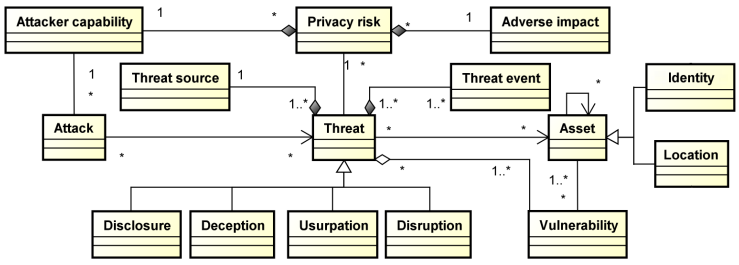

threats may cause on victims, also referred to as *adverse impact*. Analogously, the definition of risk could be easily adapted to the privacy domain by refining the concept of potential occurrence as *attacker capability*, which refers to the efforts and resources employed by the *threat source* to deploy *threat events* affecting a *privacy asset*. In the case of LBS, assets are mainly intangible. Among them, *identity* and *location* are the most important. For instance, the spatio-temporal data of an individual can be used to infer his movements and habits, to learn information about his centre of interests or even to detect a change from his usual behaviour. From these basic assets it is possible to obtain derived ones. It is to note that the asset value strongly depends on its precision. Fine-grained information concerning the who-where-when tuple is much more valuable than a coarse-grained one. For example, the value of knowing that a person, identified by the full name, will be at a place with a geo-temporal error of meters and seconds is higher than knowing that a person, only identified by the surname, will be potentially at a given city within an error window of two weeks.

In the domain of privacy, the notion of *threat* particularly refers to the *disclosure* of assets. In most of the cases, conversely to traditional security, threats are based on the potential correlation of partial informations (premises), assumed to be true, to derive more valuable information. To obtain such premises, attackers may appeal to the *usurpation* of identity, the *deception* of information or the *disruption* of services. A threat presents certain precision requirements. If such requirements are compatible with the precision provided by the asset, then the threat will be potentially feasible. For example, if a threat source requires knowing the exact current position of a user, and the only information available reveals only that the user is not at home, the threat will be hardly realisable. A threat, to be realisable, needs to exploit a *vulnerability* associated to the asset. Vulnerabilities refer to inherent weaknesses in the security procedures or internal controls concerning the asset. If there is a vulnerability without a corresponding threat, or a threat without a corresponding vulnerability, there is no risk. Next sections exploit this meta-model to approach privacy risk assessment in practice.

3.2 Risk identification

Our methodology aims at covering the existing gap in privacy risk assessment by providing a simple strategy to identify such entities in the domain of LBS following a natural from-asset-to-adverse-impact approach. The goal of risk identification is to collect the information concerning the meta-model in Figure 2 in a simple but structured way, for instance, as shown in Table 3. First, it is necessary to identify the privacy assets. Among all the *Personal Identifiable Information*, learning the location of an individual is one of the greatest threat against his privacy. Essentially because starting with the location, many other private information can be derived [4]. For example, by composing the assets in Figure 2 it is possible to derive additional ones such as the following:

– *Social relations between individuals* by considering for instance that two individuals that are in contact during a non-negligible amount of time share some

kind of social link. This information can also be derived from mobility traces by observing that certain users are in the vicinity of others.
- *Itinerary*, can be defined as a collection of locations detailed for a journey, especially a list of places to visit. An itinerary could be referred to as both physical and symbolic information containing a starting and ending point.
- *Important places*, called *Points Of Interests* (POIs), which characterise the interests of an individual. A POI may be for instance the home or place of work of an individual. Revealing the POIs of a particular individual is likely to cause a privacy breach as this data may be used to infer sensitive information such as hobbies, religious beliefs, political preferences or even potential diseases.
- *Mobility patters of an individual* such as his past, current and future locations. From the movement patterns, it is possible to deduce other informations such as the mode of transport, the age or even the lifestyle.
- *Mobility semantics* of the mobility behaviour of an individual from the knowledge of his POIs and mobility patterns. For instance, some mobility models such as *semantic trajectories* [4] do not only represent the evolution of the movements over time but also attach a semantic label to the places visited.
- *Linkable records of the same individual*, which can be contained in different geo-located datasets or in the same dataset, either anonymised or under different pseudonyms. For example, the association of the movements of Alice's car (contained for instance in dataset A) with the tracking of her cell phone locations (recorded in dataset B).

Thus, once assets fixed, reasoning about the pertinent threats related to a particular asset is easier, for example inference attacks to disclose the POIs of an individual. Then, vulnerabilities will be defined as the lack of control mechanisms enabling the potential realisation of such threats on identified assets. Finally, the users of the methodology should qualitatively determine the adverse impact on the privacy of the asset. Such an intuitive technique will guide a more systematic identification of risks on the system. An example of this technique will be detailed in Section 4.1 when presenting the case study. The rationale applied by this strategy flows in an intuitive way. We defend the simplicity of this approach as one of its major benefits. It enables even non-skilled users to handle the complexity of identifying privacy risk issues, thus easing the rest of the risk assessment.

3.3 Risk analysis

The risk analysis stage aims at estimating the privacy risk. The first step towards this goal involves proposing mechanisms to quantitatively estimate the adverse impact and the attacker capability.

Conversely to security, privacy presents a subjective component. Indeed, the adverse impact (AI) is a metric that depends on how a user perceives and interprets the importance of an asset [6]. The use of questionnaires can be very useful to estimate the adverse impact of a threatened asset. Table 1a, provides a scale to guide the intuition of users. This scale is based on four different dimensions: (i) the geographic or temporal accuracy of the asset, that can be low if it

is coarse-grained (e.g., measured in weeks or kms), or high, if it is fine-grained (e.g., measured in hours or meters); (ii) the linkability between assets, that can be low, if it is complex to relate one information with others (e.g., the license plate with the religious beliefs) or high, if it is easy to do so (e.g., the license plate with the brand of the car); (iii) the persistence of the impact, that can be low, in case the duration of the impact is transient (e.g., disclosure of the current location), or high, if it is permanent (e.g., disclosure of home address); and finally (iv) the dissemination of the assets, that can be low, if the information exposition is limited (assets are revealed to a small set of people), or high if it is publicly exposed (e.g., on the Internet).

Table 1. Adverse impact (AI) scale (a) and attacker capability (AC) scale (b). These scales are adapted to privacy from [11]. The quantitative value assigned depends on the amount of high and low scores obtained. E.g., 4 low scores and no high score involve a very low impact (rated as 0).

(a)

Accuracy	Linkability	Persistence	Dissemination	AI
Low	Low	Low	Low	0
			High	2
		High	Low	2
			High	5
	High	Low	Low	2
			High	5
		High	Low	5
			High	8
High	Low	Low	Low	2
			High	5
		High	Low	5
			High	8
	High	Low	Low	5
			High	8
		High	Low	8
			High	10

(b)

Resources requirements	Complexity	Spatial constraints	Time	AC
Low	Low	Low	Low	0
			High	2
		High	Low	2
			High	5
	High	Low	Low	2
			High	5
		High	Low	5
			High	8
High	Low	Low	Low	2
			High	5
		High	Low	5
			High	8
	High	Low	Low	5
			High	8
		High	Low	8
			High	10

Interpretation rules
4 Low + 0 High → 0 (Very low AI / AC)
3 Low + 1 High → 2 (Low AI / AC)
2 Low + 2 High → 5 (Moderate AI / AC)
1 Low + 3 High → 8 (High AI / AC)
0 Low + 4 High → 10 (Very high AI / AC)

The estimation of the attacker capability (AC) is approached by most risk assessment standards and guidelines [7, 11] through the assignation of a numeric value. While this approach is valid to provide a coarse-grained viewpoint of the threat, no detail about how threats are exploited is provided. To cover this lack, our methodology uses the concept of *attack tree*. Attack trees are structured models representing the ways in which threats are instantiated in a system [12]. In our case, the top event of attack tree represents the adverse impact achieved by attacker. To reach such an impact, the tree is structured in branches. They are defined in terms of either logic AND or OR gates, depending if the parent node's threat requires all (AND branch) or just any of them to succeed (OR branch). The branches can be considered sub-goals of the attack and can be refined until all the events become atomic actions, also referred to as leaf nodes.

Figure 3a presents an example of an attack tree to disclose who are the nearest friends of a given user. In this example, an attacker must gain access to the user profile, and then filter those friends in a radius of, e.g., 200 m. Each minimal and independent combination of leaf nodes is known as an attack scenario. Attack scenarios enable to characterise the risk from a fine-grained viewpoint, thus easily identifying privacy bottlenecks. In the previous example there are two possible attack scenarios since there are two alternatives to access to the user profile: (i)

Fig. 3. Example of attack tree and attack scenarios (a), and privacy risks levels (b).

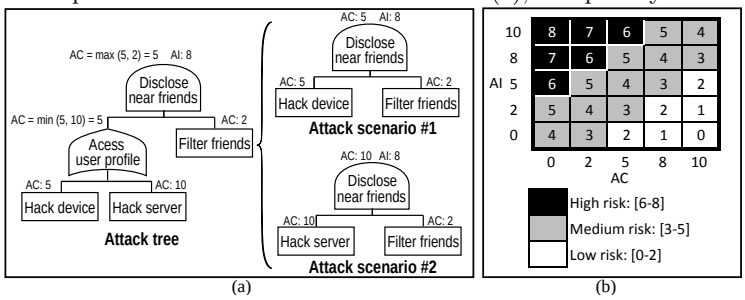

by hacking his device, or (ii) by hacking the server offering the service. Once the attack tree built, the association of quantitative values to the tree nodes can greatly improve the conclusions drawn. Since all the attack scenarios pursue the very same goal, we could map the value assigned to the adverse impact to the top event in all of them, thus keeping the compatibility with traditional standards and guidelines in the domain of security risk. Conversely, the attacker capability depends on the tasks performed in each attack scenario. Thus, the attacker capability must necessarily be computed for each leaf node of the attack tree, and then propagated towards the top event. The task requiring the maximum attacker capability typically conditions the propagation in AND branches, while OR gates propagate the minimum value.

Table 1b provides scales to normalise the attacker capability. This scale is based on four dimensions: (i) the resources required by the attacker, that can be low if the computational power is reduced (e.g., a laptop is enough), or high, if much computational power is needed; (ii) the exploit complexity required by the attacker, that can be low, if the attack is easy to launch (e.g., a script is available on the Internet) or high, if a profound knowledge is necessary to execute the attack; (iii) the spatial constraints of the attack, that can be low, if it can be performed remotely, or high, if it requires the attacker being located at some specific location (e.g., in the radio range of the victim nodes); and (iv) the duration of the necessary observation of victim nodes, that can be low if a single observation is enough, or high, if it requires many observations (e.g., for a month or even more).

We have applied the scales proposed in Figure 1 in our example in Figure 3a. As we assume the use of strong protection mechanisms to remotely access the server, all the dimensions of AC present high levels. So we have estimated as 10 the capability required by the attacker to access the server. Conversely, accessing the device as been rated as 5 because two dimensions are low: complexity, as the attacker could have downloaded an automatic malware from the Internet; and resources requirements, as this malware does not require heavy computation. As these values are the maximum values of their scenarios, they are propagated to the top event. In both cases, the adverse impact has been rated as 8, since it compromises the privacy of user's social interaction.

Following the estimation of AC and AI, risk is usually quantified using a simple expression such as Risk = AC × AI. However, such a compact and simple expression of risk implies poor semantics, which may lead to incorrect conclu-

sions. Indeed, a simple threat causing a huge damage would deliver a similar risk as a complex threat causing a negligible impact. So, we prefer to estimate risk as a tuple (AC, AI), as presented in Figure 3b. In this matrix, each cell characterises one particular risk value, in such a way that, a set of cells could be mapped to a risk level. For the purpose of this paper we have used a 0-to-8 scale to each cell considering 3 basic privacy risk levels. Thus, low, medium and high risk gathers those arrays rated from 0 to 2, from 3 to 5 and from 6 to 8 respectively. Figure 3b shows these risk levels. According to such matrix, the privacy risk associated to our example, $(5, 8)$, would be rated as 5 (medium risk).

3.4 Risk evaluation

The risk evaluation stage is in charge of ranking privacy risks regarding specific criteria. Applying different criteria, such as the importance of the asset for the business, the reputation of users or the regulation fulfilment, may lead to different rankings. According to the selected criterion, risks, regardless if they were high, medium or low, will be characterised as acceptable, tolerable and non-acceptable following an As-Low-As-Reasonably-Practicable (ALARP) [9] strategy. The risks characterised as non-acceptable will be prioritised for their treatment. However, this stage is out the scope of this paper.

4 Case study: a dynamic carpooling application

This section shows the complete cycle of our privacy risk assessment methodology through a dynamic carpooling case study. Carpooling is an urban service in which drivers and passengers share a private car from sub-urban to urban journeys or intra-urban journeys to save money while reducing the environmental pollution. Dynamic carpooling is a wireless infrastructureless version of carpooling that implements algorithms for dynamic vehicle routing problems taking into account the mobility of users. These protocols rely on intermediate users to forward carpooling requests and responses between distant users beyond one hop. Dynamic carpooling is characterised by two actors, the driver and the passenger.

Fig. 4. Sequence diagram showing the general case of dynamic carpooling.

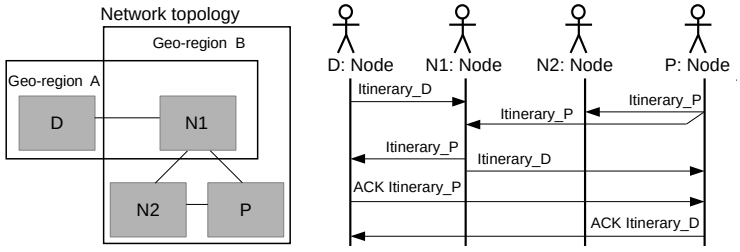

Diagram in Figure 4 illustrates the exchange of packets between Driver D and Passenger P and how these users are distributed in the network. Let us assume D belongs to the area, or geo-region A, whereas P is located at geo-region B. D and P may launch an itinerary request in which they announce the

origin GPS coordinates of the journey and the expected destination, as well as their nicknames and their preferences for the trip (e.g., if they prefer travelling with smokers or non-smokers, the accepted deviation in kms from their origin or destination, and so on). Such requests will be received by all the users located in the same geo-region. In this case, $N1$, which belongs to geo-regions A and B and thus receives the requests of D and P, is in charge of processing a potential matching between these users. If there is a matching, then $N1$ forwards requests towards D and P. If D and P accept the itinerary proposed, they will send an unicast itinerary acknowledgement to the other party.

4.1 Risk identification

The personal information considered in this case study is of prime importance since it concerns the mobility and location of users. In this section we have identified those (primary and derived) assets potentially inferred from network packets, which are of special interest given the wireless (and thus open) nature of the communication channel. These assets have been listed in Table 2.

Table 2. Assets in the dynamic carpooling application.

Primary assets	**Description**
Identity	The user nickname identifying the carpooling packet creator and the device id related to their network identifiers (IP and MAC address).
Location	The location where a carpooling user was, is or will be.
Derived assets	**Description**
Personal preferences	Details about the type of mate desired to share a trip with.
POI	Place with a special interest for the user, such as his work or home address.
Itinerary	Trip proposed and announced through carpooling packets by a driver or a passenger to go from an origin location to a destination location.
Interaction	Co-location in the domain of carpooling. It refers to the the itinerary shared by a driver and a passenger after a matching occurs.
Mobility semantics	Habits inferred from the user (e.g. the days he goes to the gym).

These assets can be exposed to multiple privacy threats. Yet, given the number of assets identified, the rest of the section will focus just on addressing some threats related to location and POI. As far as carpooling packets are exchanged through intermediate nodes using the wireless medium, malicious adversaries in the vicinity may easily capture and analyse their content. Without the aim of begin exhaustive, the set of threats introduced hereafter relies on this principle. As these threats can be combined in order to create more sophisticated threats, they will be instantiated from the simplest to the most complex attack. Table 3 synthesises the risk identification process.

1. **Eavesdropping-based location disclosure attack**: A malicious user could capture carpooling packets and analyse their content, which includes the current location of users. Since messages are forwarded, the malicious user could eavesdrop messages from distant nodes. The absence of *confidentiality* mechanisms to protect both the exact *location* and the *IP address of the device* may lead adversaries to *disclose* and match the *personal identity* with the current or future *location* of users.
2. **Triangulation-based location disclosure attack**: The complexity is higher than the eavesdropping version, since at least 3 attackers are required

to estimate the location of the victim [5]. Attackers are able to infer the location of the victim after estimating and correlating one another the signal strength and the *IP address of the device* with which carpooling packets were sent. This fact makes the attack only possible to target the nodes in the vicinity of the attackers, providing an approximate version of the user's location. However, this attack is realisable even despite the use of cryptographic mechanisms to protect the *confidentiality* packet content.

3. **POI inference attack**: Adversaries may store user's *mobility traces* to determine his POI. Mobility traces can be obtained from the instantiation of location disclosure attacks. The lack of *obfuscation* mechanisms make possible the continuous monitoring of location, thus disclosing user's POI.
4. **Fake trip negotiation attack**: Adversaries may *deceive* users requesting partners to go to a fake location. This attack may pursue economic purposes, e.g., shopping mall may want to attract users to strategically increase their sales). The attacker may exploit his previous knowledge about the victim's POI to increase the probability of success . The lack of *authentication* mechanisms to protect the *location* indicated by the application may allow adversaries to compromise users' *itinerary* and *disrupting* other potential *interactions* with legitimate drivers and passengers. Consequently, they might *influence on users' decisions* about their trip.

Table 3. Risk identification in the carpooling application.

Assets	Threats					Vulnerabilities	Description of the adverse impact
	Type				Instantiation		
	Disclosure	Usurpation	Deception	Disruption			
Location	X				#1.- Eavesdropping-based location disclosure attack	Lack of IP address confidentiality Lack of location confidentiality	Disclosure of precise user's location
Location	X				#2.- Triangulation-based location disclosure attack	Lack of IP address confidentiality	Disclosure of approximate user's location
POI		X			#3.- POI inference attack	Lack of IP address confidentiality Lack of obfuscation mechanisms	Disclosure of POI location
POI			X	X	#4.- Fake trip negotiation attack	Lack of messages authentication	Influence on user's location Reduction of availability of the service

4.2 Risk analysis

As the adverse impact finally depends on the user viewpoint, we asked potential users of carpooling[1] about their effect on their private life. According to the results of the questionnaire, scaled regarding Table 1, Attack #4 was the most concerned ($AI = 10$), followed by Attacks #1 and #3 ($AI = 8$) and Attack #2 ($AI = 5$). To determine the attacker capability we relied in our expertise. Figure 5 presents the attack trees of the attacks showed in Table 3. After applying the attacker capability propagation method presented in Section 3.2, we obtained that Attack #2 was the most complex to carry out ($AC = 8$), followed by Attacks #3 and #4 ($AC = 5$) and Attack #1 ($AC = 2$). Finally, the privacy risk, applying the matrix in Figure 3b, stated the following results. Attacks #1 and #4 presented a high privacy risk (6 points) whereas Attacks #3 and #2 presented a medium risk (with 5 and 3 points respectively).

[1] 64 people were randomly selected at LAAS-CNRS to be subjected to the experiment.

Fig. 5. Attack trees. Bold lines illustrate the most vulnerable attack scenario in case of considering attack trees with OR gates.

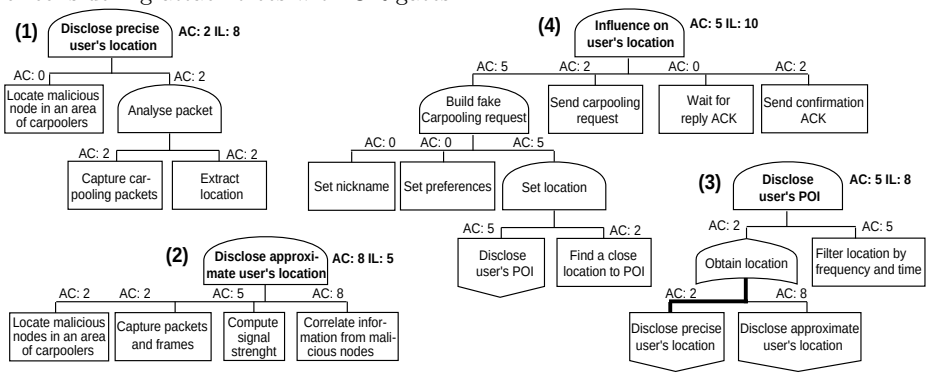

4.3 Risk evaluation

Privacy risk evaluation depends on the subjective criteria considered by the users of the methodology. To illustrate how our methodology can be customised to handle this problem, we have considered two opposite criteria. On one hand, the matrix shown in Figure 6a presents a privacy-conservative viewpoint, which could be adopted by the economic responsible of the organisation, aimed at limiting the economic investment on PETs. By playing this role, even high risks (rated up to 6 points in Figure 3b) would be considered as tolerable. In the case of dynamic carpooling, all. On the other hand, the matrix shown in Figure 6b shows a privacy-strict viewpoint, which could be adopted by the security manager, who is aware of the importance of privacy for the success of the application. Regarding this criterion, even medium risks (rated from 5 points in Figure 3b) are non-acceptable, and should be prioritised for their consequent treatment. This would be the case of the risks associated to attacks #1 (R1), #4 (R4) and #3 (R3). We claim that the adoption of the second viewpoint is better to safeguard the privacy of people and saves money to the organisation in a mid-term period.

Fig. 6. Privacy-conservative (a) and privacy-strict evaluations (b).

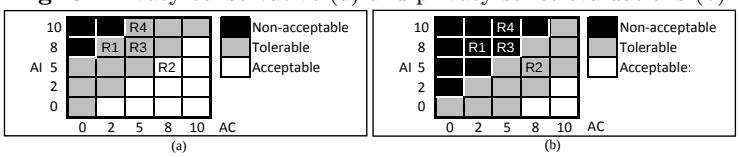

5 Conclusions

Privacy may be the greatest barrier to the long-term success of mobiquitous systems. However, despite many standards and approaches have been proposed to handle the problem of risk assessment, none, to the best of our knowledge has addressed the problem of managing the privacy risk for LBS. One of the major problems found in this paper concerns the identification of adequate information to carry out the risk assessment, as well as the way to process it. Since a lack of information may lead to obtain biased conclusions, an excess may obfuscate

the decision-making process. The asset-driven strategy proposed in Section 3 provides sufficient expressiveness to guide the rest of the privacy risk assessment. Furthermore, the risk analysis stage presented in our methodology is guided through the use of attack trees, well-known tools, to model the privacy threats, which have been expressively improved to quantify privacy risks.

The novel framework presented in this paper has been used to identify sources of risk in the lifecycle of mobiquitous solutions and find the most adequate risk trade-off between usability and privacy. Beyond this work, we are interested in studying the usefulness of our methodology to (i) guide the design PETs following a privacy-by-design approach, and (ii) compare and select (benchmark) the PETs that address the best the privacy requirements of mobiquitous systems.

Acknowledgements

This work is partially supported by the ANR French project AMORES (ANR-11-INSE-010) and the Intel Doctoral Student Honour Programme 2012.

References

[1] Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports*, 3:–, 2013.

[2] J. Dobson and P. Fisher. Geoslavery. *Technology and Society Magazine, IEEE*, 22(1):47–52, 2003.

[3] European Commission. Proposal for a regulation of the european parliament and of the council on the protection of individuals., 2012.

[4] S. Gambs, M.-O. Killijian, and M. Núñez del Prado Cortez. Show me how you move and I will tell you who you are. *Trans. on Data Privacy*, 4(2):103–126, 2011.

[5] Y. Gwon, R. Jain, and T. Kawahara. Robust indoor location estimation of stationary and mobile users. In *INFOCOM 2004. Twenty-third Annual Conference of the IEEE Computer and Communications Societies*, pages 1032–1043 vol.2, 2004.

[6] J. I. Hong, J. D. Ng, S. Lederer, and J. A. Landay. Privacy risk models for designing privacy-sensitive ubiquitous computing systems. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques*, pages 91–100. ACM, 2004.

[7] ISO27005. Information technology - security techniques - information security risk management. International Standard Organisation, 2008.

[8] ISO/IEC-Guide73. Risk management - Vocabulary - Guidelines for use in standards. International Organization for Standardization, 2009.

[9] R. E. Melchers. On the ALARP approach to risk management. *Reliability Engineering & System Safety*, 71(2):201–208, 2001.

[10] M. F. Mokbel, C.-Y. Chow, and W. G. Aref. The new casper: query processing for location services without compromising privacy. In *Proceedings of the 32nd international conference on Very large data bases*, VLDB '06, pages 763–774. VLDB Endowment, 2006.

[11] NIST800-30. Information security, guide for conducting risk assessments. U.S. Department of Commerce, NIST), 2011.

[12] I. Ray and N. Poolsapassit. Using attack trees to identify malicious attacks from authorized insiders. In *Computer Security–ESORICS 2005*, pages 231–246. Springer, 2005.