Privacy Challenges in Ambient Intelligent Systems: A Critical Discussion

Vasilis Efthymiou and **Patrice Caire** [1]

Abstract.

From the "right to be left alone" to a "commodity" that can be traded and exchanged, privacy has been defined many ways over the ages. However, there is still no consensus on one definition. Ambient Intelligence (AmI) systems power context aware, personalized, adaptive and anticipatory services. In such systems, privacy plays a critical role. The human-centered quality of AmI systems has thus prompted the use of a new kind of technology, called Privacy Enhancing Technologies (PET). Furthermore, it has now been propose to include privacy at the onset of such system design. In this survey paper, we raise the question of which specific privacy issues are raised in AmI environments and how they are addressed. We use a literature review in the fields of law, ethics, social sciences and computer sciences. We then proceed with critical discussions. We illustrate our research with a use case from Luxembourg HotCity.

1 Introduction

Many definitions of privacy are proposed, none universally accepted. Privacy is assumed to be a culturally and species-related right [36]. Meanwhile, the debate about privacy issues is steadily growing, partly due to information technology [17]. Privacy Enhancing Technologies (PET) are being developed by academia and industry [44]. Moreover, with privacy by design, privacy requirements should be taken into account at the early stage of a system design, as they potentially impact on the overall system architecture [35]. Additionally, the increasing pervasiveness of technology into our everyday lives threatens with a potential dependance on Ambient Intelligence (AmI) smart systems.

AmI, also referred to as Ubiquitous Computing and Pervasive Computing, is the Artificial Intelligence field focused on modeling, processing and even altering the context of a so-called "smart" space. The definition of this context is fundamental to the AmI system. Generally a context includes any available knowledge that can be used to describe the current environment of the system. Privacy is influenced by context [9, 19]. There are situational aspects of AmI environments that trigger different privacy concerns for different people. The emphasis is on human factors, and since AmI systems focus on assisting humans in their everyday life, privacy concerns have to be taken into consideration at the onset on system design.

AmI is set apart from other computer science domains by six specific properties: ubiquity, invisibility, sensing, memory amplification, profiling and connectedness [31, 7]. Furthermore, six basic principles, based on a set of fair information practices common in most privacy legislation, have been defined by Langheinrich [31] to guide the design of AmI systems. They are: *Notice*, i.e., users should be aware of what data are collected about them, *Choice and Consent*, *Anonymity and Pseudonymity*, *Proximity and Locality*, *Adequate Security*, and *Access and Recourse*. Wang and Kobsa [44] extend Langheinrich's principles with the 23 most frequently addressed principles in privacy laws and regulations, among which, access/participation, anonymity and choice/consent. The authors list is prompted by the observation that privacy-protecting laws exist in more than 40 countries, and typically viewing privacy from different perspectives. For example, in the US privacy is mostly self-regulated, whereas in the EU privacy is considered as a human right.

To encode privacy policies into human- and machine-readable formats, the World Wide Web Consortium (W3C) developed the Platform for Privacy Preferences Project (P3P) [2]. Unfortunately, there has been, so far, insufficient support from current web browser implementers. Similar approaches, like SPARCLE [27], or XACML [18], still have to demonstrate the practicality of their solutions. Their ambiguity, both for users and software, and lack of expressiveness keep many privacy regulations out of their scope. A shortcoming of these methods is the purely declarative nature of their policy language [1].

Currently, major threats to privacy come from personal data aggregation and the increasing power of data mining. The magnitude of the information sources and the potential to combine these sources to create a person's profile threaten individual privacy [28].

In this paper, the question we address is which specific privacy issues are raised in AmI environments and what are, currently, the proposed solutions to preserve privacy when it is threatened. This breaks down into:1) what are the definitions of privacy? 2) Which threats are specific to AmI? and 3) How is privacy preserved?

We use a literature review in the areas of ethics, law and computer science, then, proceed with critical discussions. We illustrate our findings with a scenario from a domain in which privacy plays a central role: Ambient Assisted Living (AAL), ICT-based solutions for the self-management of senior citizens in daily life activities at home. AAL addresses the problem of an aging population, which has become a main concern for many countries to insure quality of life and medical care, security and conviviality [8] to their citizens.

In this work we do not provide analytical solutions to privacy threats, nor do we cover every possible aspects of privacy. In particular, we do not include privacy in voting [22], encryption [3] nor physical approaches to privacy preservation, such as the Faraday cage approach [25] for RFID privacy issues, or privacy at the level of Wireless Sensor Networks [32].

The layout of this paper is as follows. We first, introduce our motivating scenario Chapter 2. We then present privacy definitions, Chapter 3, privacy issues in Chapter 4, and most common privacy preserving approaches in Chapter 5. We conclude in Chapter 6.

[1] University of Luxembourg, email: firstname.lastname@uni.lu

[2] http://www.w3.org/P3P/

2 MOTIVATING SCENARIO

Frank is a 70-year-old Alzheimer patient, who lives alone. His daughter, Jane, lives just a few blocks away. Usually, Frank visits Jane once or twice a week. Due to his condition, Frank has installed a Home Care System (HCS) in case he finds himself in a critical situation, and to urgently notify Jane or his friends. He also wears a health-bracelet, measuring his heart-beat, body temperature, and daily distance covered. The bracelet is connected to his smartphone, which also has a GPS and a HCS application installed. The HCS application can send vital information, such as bracelet data and current location to Frank's HCS. The HCS has a record with Frank's profile, such as name, age, address, and medical profile, as well as a list of contacts for emergency notifications. Finally, Frank carries an RFID card to verify his location. For example, Frank and Jane both have an RFID reader inside their houses: whenever Frank is near one of these readers, his location is verified.

Today, Frank is visiting Jane. He leaves his home (Figure 1, state 1) and walks to Jane's house. Suddenly, he realizes that he has been wandering about and is lost (Figure 1, state 2). He is becoming anxious. His heart is beating faster. He is sweating. Frank presses the emergency button on his bracelet and an alarm is sent to the HCS via his HCS smartphone application. The HCS infers that Frank is lost: he has been away from home for too long and has not yet checked in Jane's house. Jane is the person from the emergency list whose address is closest to Frank's current location. Thus the HCS notifies Jane about Frank's current location. The HCS also sets up Frank's smartphone voice navigator application to guide him to Jane. If Jane does not respond within five minutes, the HCS notifies the local hospital about the situation, providing Frank's medical file and current location. Finally, Jane found Frank and led him to her house with safety (Figure 1, state 3). Figure 2 depicts the connected devices.

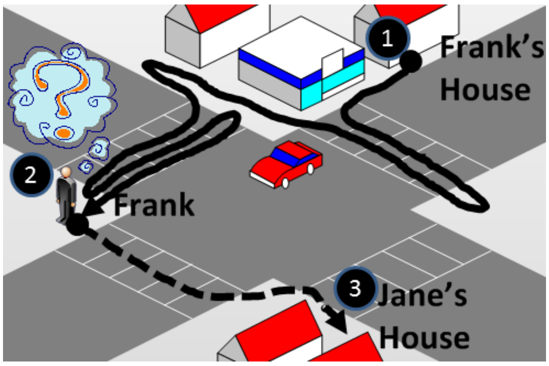

Figure 1. Frank, an Alzheimer patient, gets lost on his way to Jane's house

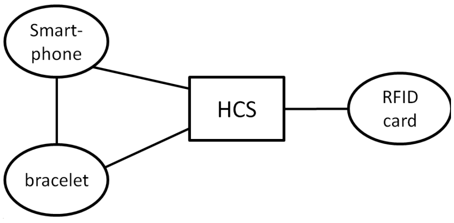

Figure 2. Connected devices

3 DEFINITIONS OF PRIVACY

The notion of privacy has been discussed extensively, not only over the last decades, but even from the 19th century. Still no globally accepted definition has been provided. Warren and Brandeis[45] are usually credited the definition of privacy as "the right to be let alone", which was actually a reference to Thomas Cooley's "Treatise on the Law of Torts"[10], written twelve years earlier, in 1878. The need for such a right emerged from the "unauthorized circulation of portraits of private persons", performed by the newspaper enterprises which used instantaneous photographs. As times changed, so did technology and its ability to intrude into people's lives. Consequently, the definition of privacy had to follow the times and incorporate these new ways of intrusion.

Some other relatively "simple" descriptions, like "exclusive access of a person to a realm of his or her own", or "control over information about oneself", even if helpful in introducing the notion of privacy, are not enough to explicitly define it. For example, there could exist many perceptions of the "realm of oneself", or the ways that someone can have "control over information". At the end, defining privacy depends on the problem of defining personal information, or even personality, notions that are mostly met on social sciences, rather than computer science. So, most definitions of privacy, if not all, take for granted (explicitly or implicitly) that these notions are well defined.

Improvements towards more specificity include Alan Westing's definition of privacy as "the ability of people to determine for themselves when, how, and to what extent information about them is communicated to others" [46] and Stefanos Gritzalis' as "the indefeasible right of an individual to control the ways in which personal information is obtained, processed, distributed, shared, and used by any other entity"[20]. Even if the latter definitions are more explicit, they still rely on the term "personal information", which can be again subjective. The Data Protection Directive defines personal data as "any information relating to an identified or identifiable natural person (the data subject)". In determining whether information concerns an identifiable person, one must apply recital 26 of the Data Protection Directive, which says that "account should be taken of all the means likely reasonably to be used either by the controller or by any other person to identify the said person". Such a definition implies a broad understanding of the notion of personal data, which may consist of all sorts of information as soon as they relate to the individual [47].

Lately, there have been so many ways in which one's privacy can be violated, that further distinction between different types of privacy needs to be made. Location privacy, for example, has been a major concern in the last few years. It can be defined, by paraphrasing Alan Westing's privacy definition, as "a special type of information privacy which concerns the claim of individuals to determine for themselves when, how, and to what extent location information about them is communicated to others" [14], or simpler as "the ability to prevent other parties from learning ones current or past location" [4]. In [39] some popular applications (Google Latitude, Facebook Places, and Glympse) are compared, with respect to their support for location privacy. Krumm finishes his survey on location privacy [29], stating that "progress in computational location privacy depends on the ability to quantify location privacy" and also noting that there was not a standard for quantifying location privacy at that time. Location-Privacy Meter, presented in [41] is an interesting tool to measure location privacy. Other types of privacy that need to be protected, as suggested in [28], include bodily privacy, territorial privacy, privacy of communications and information privacy.

Other approaches to the definition of privacy also include the idea of free, uninfluenced decision making about one's self. Kupfer [30] states that "privacy enables control over personal information as well as control over our bodies and personal choices for our concept of self", making privacy subjective to every person's own "concept of self". DeCew [13] suggests that privacy is a cluster concept covering interests in $i)$ control over information about oneself, $ii)$ control over access to oneself, both physical and mental, and $iii)$ control over one's ability to make important decisions about family and lifestyle in order to be self-expressive and to develop varied relationships.

Privacy seems to be a culturally relative right, but this doesn't mean that it is completely subjective [36]. For instance, privacy is considered to be a commodity in the US (since it relies on self-regulation) whereas in the EU it is a human right. To bridge this gap and allow US companies do business in the EU and conform to the EU Privacy Directive, EU and US arrived at an agreement, known as the safe harbor agreement. This agreements offers a convenient way of complying with the adequacy requirements of the EU Directive [42]. Privacy has nowadays become a commodity, in the sense that the consumer makes a non-monetary exchange of their personal information for value such as higher quality service and personalized offers or discounts [11].

Table 1 presents our summary categorization of the different approaches to defining privacy that have been discussed in this section. Privacy has been defined as a right, an elabler to controling personal data and, lately, as a commodity expressed through privacy policies of commercial products. This categorization is not in a chronological order, in the sense that privacy was seen as a right in the 19th century, but it can be still defined as such. However, it can be seen as a chronological categorization of the value of privacy, as it has been understood or used. It was introduced as an aspect of personal liberty, then it became an ability to control personal information and lately it has been used as a way to exchange personal information for a better service, or other commercial offers.

Right	Enabler	Commodity
-to be let alone [45] -to control the ways in which personal information is obtained, processed, distributed, shared, and used [20]	-to determine when, how, and to what extent personal information is communicated to others [46] -control over personal information as well as over our bodies and personal choices [30]	-a non-monetary exchange of consumers' personal information for value such as higher quality service and personalized offers or discounts [11].

Table 1. Categorization of privacy definitions

4 PRIVACY ISSUES IN AMI

Here we discuss the privacy problems that our scenario entails, as well as more privacy issues, typically encountered in AmI systems.

4.1 Geolocation

Geolocation can be used to breach the location privacy of a person. It is usually a problem, when past locations of a person are stored. As illustrated in [4], people probably do not care if someone finds out where they were a week ago at a specific time, but if someone could inspect the history of all their past movements, then they might start to see things differently. However, this is not always the case, since, even a single record of someone's location at a specific time can cause privacy concerns. For example, if someone is spotted in a cancer clinic, or in the office of Alcoholics Anonymous, or in a police department, then privacy could also be breached, since logical (probabilistic) assumptions could be made about this person.

In our scenario, Frank's location can be identified either by exploiting RFID privacy issues, or by acquiring, legally or not, the data stored by the HCS. For example, someone who knows that Frank lives alone and that he is currently far from home, could easily break into Frank's house, or even physically attack and rob him. Less dangerous privacy breaches include knowing Frank's wherabouts and habits, for marketing purposes, or even for surveilance reasons.

4.2 RFID Privacy Issues

RFID cards are widely used today in electronic passports, bus tickets, employee access cards, toll roads/parking access etc.; practically, on anything that needs to be identified. In some cases, RFID cards carry vital information, as in the case of electronic passports, such as name, age, address, marital status, signature, id photo etc. Other times they just carry an identification number.

In his survey[23], Juels analyzes how RFID raises two main privacy concerns: clandestine tracking and inventorying. "RFID tags respond to reader interrogation without alerting their owners or bearers. Thus, where read range permits, clandestine scanning of tags is a plausible threat." When the RFID tag also has information about the manufacturer, or the cardholder, then they are subject to inventorying. For example, an adversary could know the contents of one's bag, the amount of money he carries, the type of medication he carries, and therefore what illness he may suffer from, where he shops, his accessory preferences etc.

In our scenario, we have considered the simplest case, in which Frank's RFID card just carries an id number. Of course, this id number can be easily connected to Frank, since he is the only one who holds this specific card. If Frank simply passes by an RFID reader, placed by an adversary, similar to the one he has in his house, or in Jane's house, then he could easily be identified at the place of this reader. If Frank's card also included personal information, as in e-passports, then all these information could be at risk [24].

4.3 Patient Privacy

Lately, there has been a significant increase in the digital medical data being recorded by healthcare organizations. "While the healthcare industry has benefited from information sharing, patients are increasingly concerned about invasion of their privacy by these practices. These growing concerns on privacy led to the Health Insurance Portability and Accountability Act (HIPAA) in 2001 and have increased compliance requirements for health-care organizations" [33].

Vital information for Frank's health is stored and exchanged by the HCS, his phone, his bracelet, Jane and the local hospital. This information can be acquired by third parties, by using data mining techniques. Typically, there are three parties involved in the privacy problem in data mining [33]: $i.)$ the data owner (the organization that owns the data) who wants to discover knowledge from the data without compromising the confidentiality of the data, $ii.)$ individuals who provide their personal information to the data owner and want their privacy protected and $iii.)$ the third party data user who has access to the data released by the data owner.

This third party can be an individual data miner (either insider or outsider to the data owner), or an organization that has a data sharing agreement with the data owner. In our example, the local hospital could be the third party, or even a medical company that has a data sharing agreement with the hospital. Even if the data sent to the third party are de-anonymized, they can be combined with publicly available data and still identify the refered individual.

Another interesting source of privacy breach iw provided by people who are authorized to access patient data. "Recent studies have revealed that numerous policy violations occur in the real world as employees access records of celebrities, family members, and neighbors motivated by general curiosity, financial gain, child custody lawsuits and other considerations" [5].

4.4 Personal Data Leackage

Personal data, namely any information relating to an identified or identifiable person, could be considered as a superset of patient data. Combinations of few characteristics can be used to uniquely or nearly uniquely identify some individuals. It is discussed in [43] that 87% of the population in the United States had reported characteristics that likely made them unique based only on ZIP code, gender and date of birth. For example, just by buying the voter registration list for Cambridge Massachusetts and having a copy of publicly available, anonymized, patient-specific data, Sweeney could identify the patient record of the governor of Massachussetts at that time. "Clearly, data released containing such information about these individuals should not be considered anonymous. Yet, health and other person-specific data are often publicly available in this form." A similar example was provided in [37], which presents a framework that analyzes privacy and anonymity in social networks and re-identifies anonymized social network graphs. A third of the users who could be verified to have accounts on both Twitter and Flickr, could be re-identified in the anonymous Twitter graph with only a 12% error rate.

In our scenario, Frank's age, address, medical profile, health data, whereabouts, marital status etc. could be available to third parties, by data aggregation, without Frank's explicit authorization.

4.5 Unauthorized Actions

As discussed in chapter 3, some definitions of privacy also include the aspect of control over personal choices. When Frank decides to push the button on his bracelet, he implicitly gives authorization to his HCS to take action. However, there could be a case, in which he does not push the button and the HCS realizes that there is an emergency. If the HCS is programmed to call for help and share Frank's medical profile, then that could be a breach of his privacy.

Even if Frank agrees to share his medical record, there is also an issue regarding the recipient of this information. Frank could accept sharing this information with the local hospital, but disapprove sharing the same information with his daughter. To avoid such kind of conflicts, authorization rules have to be predefined by Frank.

4.6 Discussion

In this section we have seen ways in which Frank's privacy could be breached. By recording Frank's past locations, an adversary could infer important information about Frank's personal life. Even by knowing Frank's current location, an adversary could physically attack him, or break into his house. If an RFID card is used, then again Frank could be spotted. Moreover, if this RFID card carries personal information, or if multiple RFID cards are used for things Frank carries with him, then these data could be at risk. Vital information about Frank's health, stored and transmitted by his devices, could be acquired by third parties without his authorization. Data aggregation could make it possible for an adversary to infer Frank's personal data, like his age, address, medical profile, marital status etc. Finally, Frank's own HCS could breach his privacy, by taking important decisions about Frank's life, without his approval.

5 PRESERVING PRIVACY

In this section we present some typical approaches to preserve privacy. Even if these approaches are not solely focusing on AmI systems, correlating them with such systems is trivial. When necessary, we use our scenario from chapter 2 to illustrate how the presented approach could be applied in AmI systems.

We have divided this section in two categories; the first one is for data that can be accessed, but not in their original form, and the second mainly focuses on techniques that will provide personal information to a specific group of people, while hide it from anyone else. There are also other approaches, like auditing mechanisms [5], which could complement privacy policies.

5.1 Modifying Available Data

Privacy issues occur when someone's personal data become available, against this person's will. However, there is no issue at all when the same personal data is available, but without the possibility, or, to be more realistic, with a very small chance of connecting them to this person. For example, it is certainly a breach of privacy to know that your neighbour, Frank, has the Alzheimer's disease, when Frank doesn't want you to know that. However, Frank would have no problem if it was publicly available that someone with the pseudonym X suffers Alzheimer's disease. The property of being indistinguishable among a set of individuals is called anonymity. The problem is that even anonymized data can be combined and finally identify who X is. There is no privacy issue, either, to know that a person with the pseudonym X, lives on 24, Monterey street. But if we know that the same person always gets the same pseudonym, then we can easily infer that the person who lives on 24, Monterey street suffers from the Alzheimer's disease and in a similar manner that his name is Frank.

A very popular approach during the last decade was the notion of k-anonymity [43]. In a k-anonymized dataset, each record is indistinguishable from at least k - 1 other records with respect to certain identifying attributes. However, a more recent work [34], introducing ℓ-diversity, has proven that k-anonymity does not guarantee privacy against attackers using background knowledge.

Dwork [15, 16] introduces, and describes a mechanism achieving the notion of differential privacy. It is based on Dalenius' [12] desideratum for statistical databases, which states that nothing about an individual should be learnable from the database that cannot be learned without access to the database.

A completely different approach is presented in [40]. The Personal Data Stream (PDS) is designed to give users new data management tools, based on three foundational design principles: primacy of participants, data legibility, and engagement of participants throughout the data life cycle. With the PDS, the participants are in control of their data, able to make privacy decisions. A prerequisite for this approach is that participants should be able to understand what the data mean and reveal about them.

5.2 Data Access: A Logical Approach

Solutions to privacy based on logic, mainly focus on permissions to access data (authorization problem). Following the definitions of privacy, it should be the persons whose data are shared that decide who will receive their personal data, either directly, e.g. by being asked, or even just aware each time their data is broadcasted, or indirectly, by agreeing upon a privacy policy. In each case, they should always be in a position to control the flow of their personal data. In our scenario, this kind of privacy preservation would include Frank participating in the design of his HCS privacy policy, by stating his privacy preferences. For example, he could state that only his personal doctor can have access to his medical profile and, in the case of an emergency, this access could be also granted to any other doctor in duty. This would prohibit Jane from viewing Frank's medical profile.

DEAL [19] is a formal high-level authorization language, aiming to specify access control policies in open and dynamic distributed systems. It supports negative authorization, rule priorities, hierarchical category authorization and nonmonotonic reasoning.

Aucher et al. [1, 2] also study how to formally specify and reason about privacy policies in terms of permitted and forbidden knowledge by using epistemic logic and deontic logic, branches of modal logic. The requirements the authors set on languages for specifying and reasoning about privacy policies are that by using such languages, one should be able to: $i.$) distinguish between a permission to know and the permission to send a message, $ii.$) specify and reason about the order in which messages can be sent, $iii.$)specify obligations in privacy policies and $iv.$) express meta-security policies.

The Coprelobri (computers and privacy regulations: the logical bridge) project [1], is built on a logical language that can be used to represent and reason on privacy policies. It could be used to provide writing assistance to lawyers in charge of making privacy policies, regulations and law. It could also be used to check that a given policy is compliant with a set of high-level regulations. Deontic Logic for Privacy (DLP logic) [38] is a normal deontic temporal language, which can represent information about personal data usage and protection. DLP can deal with deontico-temporal notions which are prominent in privacy-related regulations.

Collaboration and privacy are two competing concepts. Kanovich et al. [26] discuss the interplay between confidentiality, or policy compliance, and goal reachability. The authors focus on the research question whether the agents can achieve their common goal while having some confidentiality guarantees. "The main confidentiality concern is that data might become available or visible to an agent who is prohibited from viewing it according to one of the policies." It is assummed that each agent has a data confidentiality policy which specifies which pieces of data other agents are prohibited from learning. Affine Logic (AL) is used to model the reachability of partial goals, because it allows working with the relevant resources in arbitrary contexts.

PROTUNE (Provisional Trust Negotiation) [6] is a system for specifying and cooperatively enforcing security and privacy policies. Protune relies on logic programming for representing policies and for reasoning with and about them. "The use of set of Horn rules for policies together with ontologies provide the advantage of well-defined semantics and machine interoperability, hence allowing for automated negotiations." In Protune, policies are basically sets of Horn rules, on which the system has to perform several kinds of symbolic manipulations such as deduction, abduction, and filtering. "Policies are monotonic in the sense that, as more credentials are released and more actions executed, the set of permissions does not decrease." Protune introduces a mechanism for answering why, why-not, how-to, and what-if queries on rule-based policies.This mechanism aims to help common users become aware of the policy applied by the systems they interact with and even take control over it.

In [21], a semantically rich, policy-based framework that constrains the information flow in a context-aware system is presented. It uses an OWL ontology to represent dynamic aspects of context-aware systems and a combination of OWL-DL and Jena rules specifying the policy to perform reasoning. It enforces user's privacy preferences using static information about the user as well as dynamic information observed and inferred from the context. "Privacy preferences are access control rules that describe how a user wants to share which information, with whom, and under what conditions." This framework provides users with appropriate levels of privacy to protect the personal information, including the possible inferences from this information, on their mobile devices.

5.3 Discussion

It should be clarified that the two approaches to preserve privacy, presented in this chapter are not used for the same purpose. Data masking/ anonymizing etc is typically used when the personal data are expected to be accessed by third parties. For example, a hospital that wants to share scientific data, based on patient records, while at the same time preserve the patients' privacy, would use one of these solutions. In other words data masking is all about **what** kind of data will a third party have access to. On the other hand, the authorization problem is about deciding **who** will have access to personal, typically not anonymized data. However, a combination of these approaches would be interesting, since we would expect to have control of **who** has access and to **what** type of data.

6 CONCLUSION

With the pervasiveness of AmI systems, privacy issues and how to preserve own's own privacy has become key. In this paper we raise the question of which specific issues and solutions are currently used in AmI environment. First, to understand how the concept of privacy has been used up to now, we present the privacy definitions put forward in the literature. We note the multi-faceted aspect of the concept, ranging from being a "right to be let alone" [10, 45], to enabling "control over personal information" [30, 13, 46], to a mere "commodity" [11]. We then, present and discuss the most common privacy issues pertaining to AmI environments. We find that geolocation could become a very serious privacy breach, especially when a history of locations is recoreded. RFID is a promising way of identification, but information stored in an RFID tag can be at risk, if RFID security is not thoroughly designed. Patient data can be acquired by third parties with data aggregation techniques, even from publicly available, anonymized data. Finally, we describe the two most common approaches used to preserve privacy, namely by modifying the available data and providing authorization mechanisms. These approaches answer to two different questions, namely **what** kind of data can a third party acquire and **who** can acquire private data respectively. However, they can be combined and enable control over both these questions. Throughout our paper, we illustrate how these issues and techniques may arise in a real-life situation with a motivating scenario validated by the HotCity of Luxembourg set in the Ambient Assisted Living (AAL) domain.

In future works, we will include privacy in the context of AAL and AmI systems and perform reasoning on the authorization problems. Additionally, in order to respect users' privacy preferences and endow the system with user-friendliness and conviviality, we will address the trade-offs that must be done for an AmI system to be both private and convivial.

REFERENCES

[1] Guillaume Aucher, Catherine Barreau-Saliou, Guido Boella, Annie Blandin-Obernesser, Sébastien Gambs, Guillaume Piolle, and Leendert

Van Der Torre, 'The Coprelobri project : the logical approach to privacy', in *2e Atelier Protection de la Vie Privée (APVP 2011)*, Sorèze, France, (June 2011).

[2] Guillaume Aucher, Guido Boella, and Leendert Van Der Torre, 'Privacy policies with modal logic: the dynamic turn', in *Proceedings of the 10th international conference on Deontic logic in computer science*, DEON'10, pp. 196–213, Berlin, Heidelberg, (2010). Springer-Verlag.

[3] Josh Benaloh, Melissa Chase, Eric Horvitz, and Kristin Lauter, 'Patient controlled encryption: ensuring privacy of electronic medical records', in *Proceedings of the 2009 ACM workshop on Cloud computing security*, CCSW '09, pp. 103–114, New York, NY, USA, (2009). ACM.

[4] Alastair R. Beresford and Frank Stajano, 'Location privacy in pervasive computing', *IEEE Pervasive Computing*, **2**, 46–55, (2003).

[5] Jeremiah Blocki, Nicolas Christin, Anupam Datta, and Arunesh Sinha, 'Regret minimizing audits: A learning-theoretic basis for privacy protection', in *Proceedings of the 2011 IEEE 24th Computer Security Foundations Symposium*, CSF '11, pp. 312–327, Washington, DC, USA, (2011). IEEE Computer Society.

[6] Piero A. Bonatti, Juri L. Coi, Daniel Olmedilla, and Luigi Sauro, 'Policy-driven negotiations and explanations: Exploiting logic-programming for trust management, privacy & security', in *Proceedings of the 24th International Conference on Logic Programming*, ICLP '08, pp. 779–784, Berlin, Heidelberg, (2008). Springer-Verlag.

[7] Philip Brey, 'Freedom and privacy in ambient intelligence', *Ethics and Information Technology*, **7**, 157–166, (2005).

[8] Patrice Caire and Leendert van der Torre, 'Convivial ambient technologies: Requirements, ontology and design', *Comput. J.*, **53**(8), 1229–1256, (2010).

[9] Diane J. Cook, Juan C. Augusto, and Vikramaditya R. Jakkula, 'Ambient intelligence: Technologies, applications, and opportunities.', *Pervasive and Mobile Computing*, 277–298, (2009).

[10] T.M.I. Cooley, *A Treatise on the Law of Torts: Or the Wrongs which Arise Independent of Contract*, Callaghan, 1878.

[11] Mary J. Culnan and Robert J. Bies, 'Consumer privacy: Balancing economic and justice considerations', *Journal of Social Issues*, **59**(2), 323–342, (2003).

[12] T. Dalenius, 'Towards a methodology for statistical disclosure control', *Statistik Tidskrift*, **15**(429-444), 2–1, (1977).

[13] J.W. DeCew, *In Pursuit of Privacy: Law, Ethics, and the Rise of Technology*, G - Reference, Information and Interdisciplinary Subjects Series, Cornell University Press, 1997.

[14] M. Duckham and L. Kulik, 'Location privacy and location-aware computing', *Dynamic & mobile GIS: investigating change in space and time*, 34–51, (2006).

[15] Cynthia Dwork, 'Differential privacy', in *ICALP (2)*, eds., Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, volume 4052 of *Lecture Notes in Computer Science*, pp. 1–12. Springer, (2006).

[16] Cynthia Dwork, 'A firm foundation for private data analysis', *Commun. ACM*, **54**(1), 86–95, (January 2011).

[17] Crystal Edge and William Hafner, 'Analysis of privacy trends over the 19th and 20th centuries', in *Proceedings of the Fourth Annual AIS SIGSEC Workshop on Information Security and Privacy (WISP 2009)*, Phoenix, AZ, USA, (2009).

[18] Organization for the Advancement of Structured Information Standards., 'Privacy policy profile of xacml v2.0', Technical report, (2005).

[19] Irini Genitsaridi, Antonis Bikakis, and Grigoris Antoniou, 'Deal: A distributed authorization language for ambient intelligence', *IJACI*, **3**(4), 9–24, (2011).

[20] Stefanos Gritzalis, 'Enhancing web privacy and anonymity in the digital era', *Inf. Manag. Comput. Security*, **12**(3), 255–287, (2004).

[21] Pramod Jagtap, Anupam Joshi, Tim Finin, and Laura Zavala, 'Preserving Privacy in Context-Aware Systems', in *Proceedings of the Fifth IEEE International Conference on Semantic Computing*. IEEE Computer Society Press, (October 2011).

[22] Hugo Jonker and Jun Pang, 'Bulletin boards in voting systems: Modelling and measuring privacy', in *ARES*, pp. 294–300. IEEE, (2011).

[23] Ari Juels, 'Rfid security and privacy: A research survey', *Journal of Selected Areas in Communication (J-SAC)*, **24**(2), 381–395, (2006).

[24] Ari Juels, David Molnar, and David Wagner, 'Security and privacy issues in e-passports', in *Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks*, SECURECOMM '05, pp. 74–88, Washington, DC, USA, (2005). IEEE Computer Society.

[25] Ari Juels, Ronald L. Rivest, and Michael Szydlo, 'The blocker tag: selective blocking of rfid tags for consumer privacy', in *Proceedings of the 10th ACM conference on Computer and communications security*, CCS '03, pp. 103–111, New York, NY, USA, (2003). ACM.

[26] Max Kanovich, Paul Rowe, and Andre Scedrov, 'Collaborative planning with confidentiality', *Journal of Automated Reasoning*, **46**, 389–421, (2011). 10.1007/s10817-010-9190-1.

[27] Günter Karjoth and Matthias Schunter, 'A privacy policy model for enterprises', in *Proceedings of the 15th IEEE workshop on Computer Security Foundations*, CSFW '02, pp. 271–, Washington, DC, USA, (2002). IEEE Computer Society.

[28] Maria Karyda, Stefanos Gritzalis, Jong H. Park, and Spyros Kokolakis, 'Privacy and fair information practices in ubiquitous environments: Research challenges and future directions', *Internet Research*, **19**(2), 194–208, (2009).

[29] John Krumm, 'A survey of computational location privacy', *Personal Ubiquitous Comput.*, **13**(6), 391–399, (August 2009).

[30] Joseph Kupfer, 'Privacy, autonomy, and self-concept', *American Philosophical Quarterly*, **24**(1), pp. 81–89, (1987).

[31] Marc Langheinrich, 'Privacy by design principles of privacy-aware ubiquitous systems', in *Ubicomp 2001: Ubiquitous Computing*, eds., Gregory Abowd, Barry Brumitt, and Steven Shafer, volume 2201 of *Lecture Notes in Computer Science*, pp. 273–291. Springer Berlin / Heidelberg, (2001).

[32] Na Li, Nan Zhang, Sajal K. Das, and Bhavani Thuraisingham, 'Privacy preservation in wireless sensor networks: A state-of-the-art survey', *Ad Hoc Networks*, **7**(8), 1501 – 1514, (2009). Privacy and Security in Wireless Sensor and Ad Hoc Networks.

[33] Xiao-Bai Li and Luvai Motiwalla, 'Protecting patient privacy with data masking', in *Proceedings of the Fourth Annual AIS SIGSEC Workshop on Information Security and Privacy (WISP 2009)*, Phoenix, AZ, USA, (2009).

[34] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubramaniam, 'L-diversity: Privacy beyond k-anonymity', *ACM Trans. Knowl. Discov. Data*, **1**(1), (March 2007).

[35] Daniel Le Métayer. Privacy by design : towards a systematic approach. Atelier Protection de la Vie Privée (APVP 2011), June 2011.

[36] Adam D. Moore, 'Privacy: Its meaning and value', *American Philosophical Quarterly*, **40**(3), pp. 215–227, (2003).

[37] Arvind Narayanan and Vitaly Shmatikov, 'De-anonymizing social networks', in *Proceedings of the 2009 30th IEEE Symposium on Security and Privacy*, SP '09, pp. 173–187, Washington, DC, USA, (2009). IEEE Computer Society.

[38] Guillaume Piolle and Yves Demazeau, 'Representing privacy regulations with deontico-temporal operators', *Web Intelligence and Agent Systems*, **9**(3), 209–226, (2011).

[39] Marcello Paolo Scipioni and Marc Langheinrich, 'Towards a new privacy-aware location sharing platform', *Journal of Internet Services and Information Security*, **1**, (2011).

[40] Katie Shilton, Jeffrey A Burke, Deborah Estrin, and Mark Hansen, 'Designing the personal data stream : Enabling participatory privacy in mobile personal sensing', in *37th Research Conference on Communication, Information and Internet Policy (TPRC)*, pp. 25–27, Arlington, VA, (September 2009).

[41] Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, and Jean-Pierre Hubaux, 'Quantifying location privacy', in *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, SP '11, pp. 247–262, Washington, DC, USA, (2011). IEEE Computer Society.

[42] Gerhard Steinke, 'Data privacy approaches from us and eu perspectives', *Telematics and Informatics*, **19**(2), 193 – 200, (2002). Regulating the Internet: EU and US perspectives.

[43] Latanya Sweeney, 'k-anonymity: A model for protecting privacy.', *International Journal of Uncertainty, Fuzziness & Knowledge-Based Systems*, **10**(5), 557, (2002).

[44] Yang Wang and Alfred Kobsa, 'Privacy-enhancing technologies', in *Handbook of Research on Social and Organizational Liabilities in Information Security*, pp. 352–375. Information Science Reference, (2008).

[45] Samuel D. Warren and Louis D. Brandeis, 'The right to privacy', *Harvard Law Review*, **4**(5), pp. 193–220, (1890).

[46] Alan Westin, *Privacy and Freedom*, New Jork Atheneum, 1967.

[47] David Wright, Serge Gutwirth, Michael Friedewald, Paul De Hert, Marc Langheinrich, and Anna Moscibroda, 'Privacy, trust and policy-making: challenges and responses', *Computer Law & Security Report*, **25**(1), (February 2009).