

Privacy Challenges in Ambient Intelligence Systems

Lessons Learned, Gaps and Perspectives from the AAL Domain and Applications

Patrice Caire^a, Assaad Moawad^a, Vasilis Efthymiou^b, Antonis Bikakis^c and Yves Le Traon^a

^a *Interdisciplinary Center for Security, Reliability and Trust (SnT),
University of Luxembourg, Luxembourg*

E-mail: first.lastname@uni.lu

^b *Foundation for Research and Technology - Hellas (FORTH)*

Institute of Computer Science

E-mail: vefthym@ics.forth.gr

^c *University College London*

Department of Information Studies

E-mail: a.bikakis@ucl.ac.uk

Abstract. Today, privacy is a key concept. It is also one which is rapidly evolving with technological advances, and there is no consensus on a single definition for it. In fact, the concept of privacy has been defined in many different ways, ranging from the “right to be left alone” to being a “commodity” that can be bought and sold. In the same time, powerful Ambient Intelligence (AmI) systems are being developed, that deploy context-aware, personalised, adaptive and anticipatory services. In such systems personal data is vastly collected, stored, and distributed, making privacy preservation a critical issue. The human-centred focus of AmI systems has prompted the introduction of new kinds of technologies, e.g. Privacy Enhancing Technologies (PET), and methodologies, e.g. Privacy by Design (PbD), whereby privacy concerns are included in the design of the system. One particular application field, where privacy preservation is of critical importance is Ambient Assisted Living (AAL). Emerging from the continuous increase of the ageing population, AAL focuses on intelligent systems of assistance for a better, healthier and safer life in their living environment. In this paper, we first build on our previous work, in which we introduced a new tripartite categorisation of privacy as a right, an enabler, and a commodity. Second, we highlight the specific privacy issues raised in AAL. Third, we review and discuss current approaches for privacy preservation. Finally, drawing on lessons learned from AAL, we provide insights on the challenges and opportunities that lie ahead. Part of our methodology is a statistical analysis performed on the IEEE publications database. We illustrate our work with AAL scenarios elaborated in cooperation with the city of Luxembourg.

Keywords:

Ambient Assisted Living (AAL), Ambient Intelligence Systems, Privacy, Privacy Enhancing Technology (PET), Privacy by Design (PbD), privacy preservation, data access control

1. Introduction

The debate about privacy issues is paramount, and this is mainly due to the impact of information technology [1]. Moreover, many different definitions of privacy have been proposed, but up to now and to

our knowledge, no single one has been universally accepted.

Currently, privacy is assumed to be a right to be preserved [2]. To enforce this right, methodologies such as *Privacy by Design* (PbD) have pushed forward privacy requirements so that they be taken into account

at the early stage of a system design. Indeed, they may impact the overall system architecture [3].

Furthermore, the increasing pervasiveness of technology in our everyday lives adds a new threat: our dependence on this very technology. It is precisely to address this threat that *Privacy Enhancing Technologies* (PET) are developed. Such developments are also part of new approaches stemming from the concept of *Ambient intelligence* (AmI).

Also referred to as *Ubiquitous Computing* and *Pervasive Computing*, AmI is the Artificial Intelligence field focused on modeling, processing and also altering the context of the so-called *smart spaces*. The definition of *contexts* is fundamental to AmI systems. Typically, a context includes any available information that can be used to describe the environment of the system as well as the system itself. Some situational aspects of AmI environments however may trigger different privacy concerns for different people, showing that privacy is influenced by contexts [4,5]. Furthermore, since AmI systems focus on assisting humans in their everyday life, and emphasising human factors, privacy concerns have to be taken into consideration at the onset of the design of such systems.

One application domain of AmI where privacy requirements are of critical importance, due in part to the potential fragility of its users, is *Ambient Assisted Living* (AAL). AAL has emerged out of the continuous increase of the older population in Europe and worldwide. This has called for new technological solutions for improving the health [6,7], quality of life [8,9] and independent living [10,11] of older citizens. AAL has hence arisen as a multi-disciplinary field to exploit information and communication technologies in healthcare systems aiming at countering the effects of a growing elderly population.

Today, AAL systems are developed for personalised, adaptive and high service quality to satisfy requirements such as privacy, interoperability, usability, security, and accuracy. Such systems need to provide services that are not only sensitive and responsive, but also unobtrusively integrated into our daily environments [12,13,14].

AAL, and more generally AmI systems, give more control to humans. In [15], Brey notices that AmI on one hand, tends to make users' environments responsive to their intended actions and on the other hand, provides customised information to ease their lives by reducing the cognitive and physical effort required to perform certain tasks. But this gain comes at a cost: "delegating control to machines". This brings us to the

following paradox: users gain control by losing control – as they give it to their personalised systems.

The question of trade-off between gaining and losing control, brought about by such systems is highlighted by Bohn et al. [16]. The authors point out that sayings such as, "the walls have ears" and "if these walls could talk..." have now become a somewhat disturbing reality. This represents obvious challenges to privacy, which need to be addressed. Recently, a framework to proactively embed privacy directly into system design, was approved as the international standard for ensuring privacy in the information era. It is called *Privacy by Design PbD* [17].

PbD has been applied in many different areas, such as Smart Grids, biometric encryption, cloud computing, surveillance systems and others¹. PbD is based on seven fundamental principles: *privacy as the default setting; end-to-end security; respect for user privacy; openness and transparency; proactive not reactive, preventative not remedial; privacy embedded into design; and full functionality - positive-sum, not zero-sum*. These principles were set to extend the PETs and were designed to "be applied with special vigour to sensitive data such as medical information and financial data" [17].

After analysing the properties of pervasive computing environments that could set apart AmI from other domains, six specific properties were highlighted by [15,18]. These are: ubiquity, invisibility, sensing, memory amplification, profiling and connectedness. Pursuing this research, and in order to guide the design of AmI systems, Langheinrich developed six principles based on a set of fair information practices common in most privacy legislation [18], namely: *notice, choice and consent, anonymity and pseudonymity, proximity and locality, adequate security, and access and recourse* [15,18].

More recently, an extension to Langheinrich's principles was proposed by Wang and Kobsa [19]. Altogether, these principles represent the 23 most frequently addressed principles in privacy laws and regulations. Particularly to be noted, are those related to access/participation, anonymity and choice/consent. The authors started their list by observing that privacy-protecting laws exist in more than 40 countries, which typically view privacy from different perspectives. For example in the US, privacy is mostly self-regulated,

¹<http://www.privacybydesign.ca>: the PbD website contains more information about successful application examples

whereas in the EU, privacy is considered a human right.

As expected, the Web is catching up with addressing privacy by moving forward with a number of approaches, being thereby a step ahead of AmI. For example, to encode privacy policies into human-readable 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 [20], or XACML [21], 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 [22].

Currently, major threats to privacy come from personal data aggregation and the increasing power and sophistication of data mining technics. The magnitude of the information sources and the increasing potential to combine these sources to for example, create a person's profile, greatly threaten individual privacy [23]. This is particularly critical in the field of AAL, in which users are often not aware of the potential mischievous use of their personal information.

All the above challenges bring about our main research question as follows:

Research question: Which are the specific privacy issues raised and addressed in the AAL domain and applications?

This breaks down into the following sub-questions:

1. What are the definitions of privacy in social science and computer science?
2. What are the privacy threats specific to the AAL domain?
3. How is privacy currently preserved in AAL applications?
4. What are the open questions and the future trends concerning privacy threats in the AAL domain and applications?

To address these questions, our methodology consists in reviewing existing literature in the areas of ethics, law and computer science. Furthermore, we conduct a statistical and quantitative analysis. We then

proceed with critical discussions. We illustrate our findings with an AAL scenario.

Typically, AAL proposes ICT-based solutions to allow senior citizens living at home to better self-manage their daily activities. According to Cook et al. [4], there is a tremendous need for research on AmI to improve the quality of life for people with disabilities and to promote aging at home. In the authors' view, AAL applications demonstrate "the tension found between two noble goals: preserving privacy and providing useful smart environment benefits."

The pressure is also real as by 2040, there will be 23% of the population over 65 years old [24] and over 11 million people will suffer from dementia related to the Alzheimer's disease [25]. This will also affect the economy, since, for example, the long-term projected total losses to the US economy is expected to be nearly two trillion dollars [26], given that the costs of US nursing home care are \$40,000 a year [27]. Hence, the AAL domain has become a main concern for many countries seeking to insure quality of life, medical care, security and conviviality [28] to their citizens.

In this research, we partnered with Luxembourg HotCity, the city-wide mesh Wi-Fi network, which has AAL expertise, to create and validate a number of use cases from which originated our motivating scenario. Initiated in 2007 by the City of Luxembourg to build Europe's most advanced municipal Wi-Fi network, HotCity³ is geared to providing seamless coverage across the city's 50 square kilometres area. Four new cities in Luxembourg and 10 applications were recently added. Today, the HotCity network serves a growing number of cities in Luxembourg and has two proofs-of-concept, namely in Belgium and Holland. HotCity reinforces the important role of a municipal Wi-Fi connectivity by contributing to bringing together better connected businesses and social communities. One of its priorities is to adequately prepare its AAL systems for its senior citizens by anticipating privacy challenges.

Regarding the scope of this article, based on our previous work [29], we do not provide solutions to privacy threats, nor do we cover every single aspect of privacy. In particular, we do not include privacy in voting [30], encryption [31] nor the physical approaches to privacy preservation, such as the Faraday cage [32] for RFID cards, or privacy for vehicular networks, or

²<http://www.w3.org/P3P/>

³<https://www.hotcity.lu/en/laptop/www>

at the level of Wireless Sensor Networks [33]. Furthermore, although our field of study is AAL, discussing in detail the many ethical and medical challenges related to patients' healthcare is beyond the scope of this article. For a thorough study of various ethical issues that arise in Intelligent Environments, such as AmI environments and AAL, as well as for the means of coping with them, we recommend reading the work performed within the eFRIEND framework [34].

The layout of this article is as follows. First, we introduce our motivating scenario in Section 2. Then, in Section 3, we present privacy definitions, introduce a new categorisation of privacy definitions, and a view on the evolution of the concept of privacy over the last forty years. In Section 4, we review main privacy issues specific to AAL, and present a statistical analysis based on the appearance of the term in the IEEE database of publications. Then, we summarise the most common privacy preserving approaches in Section 5. In Section 6, we discuss the lessons learned throughout this research, and present our insights on what lies ahead concerning privacy in AAL and AmI environments. Finally, we conclude in Section 7, and present in the Appendix, the supporting material for our references selection process.

2. MOTIVATING SCENARIO

An important goal for our HotCity partner is to encourage social inclusion by providing Internet and application access to all citizens, and to foster communication and interaction among the wide range of people in Luxembourg, particularly the ones who may be excluded from the social dialogue. Hence, HotCity focuses on senior citizens, who, for example, may need continuous monitoring and particularly cares to enhance their lives in the city.

Leveraging HotCity's interest, expertise and knowledge in the AAL domain, we elaborated twelve use case scenarios. A complete description and analysis of these scenarios can be found in [35]. The selection of the scenarios was done by the HotCity experts who ranked each scenario. We now present the scenario selected as our running example.⁴

⁴The ranking was based on the following two criteria: (1) likelihood, i.e., the probability that the scenario occurs and (2) impact, i.e., the consequence on human life of the failure of the scenario. As in risk based testing approaches [36], likelihood and impact have been used to prioritize scenarios, from low (value 1) to

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 of Frank's profile, which includes his 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 close to 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 Electronic Patient Record (EPR) 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.

In this scenario, Frank's privacy could be breached in multiple ways. For example, when the HCS notifies the local hospital about Frank's situation, it is without specifically asking Frank at that very moment. Indeed, Frank could expect that only his daughter is notified, as he did not believe that his situation was so critical

high (value 3). The *priority* P of each scenario is calculated as the product of *likelihood* L and *impact* I , i.e. $P = L \times I$. The results describe the relevancy of these twelve scenarios. Further analysis can check whether specific requirements, e.g. conviviality, user-friendliness and security, are satisfied.




Fig. 1. Frank gets lost on his way to Jane's house

to require notifying the hospital. Also, someone could easily identify Frank's whereabouts, by setting RFID readers across the city, or by receiving the transmissions of Frank's smartphone. This could expose him to have his house broken into when he is away, or when he is visiting a hospital for some days. Moreover, creating Frank's profile in order to improve his everyday life, could work against him.

The designers of AmI systems therefore must be aware of the privacy challenges that arise with such systems. They also must be able to prevent them. Following the privacy by design principles [18], system designers will have to set up the default configuration of the system to the maximum privacy first. When installing the system, Frank (the end user) can select the choices that suit him well, thus giving his apriori consent to the potential privacy trade-off.




Fig. 2. Connected devices

3. DEFINITIONS OF PRIVACY

The notion of privacy has been discussed extensively, not only over the last decades, but even since the 19th century, and appears in the literature of var-

ious disciplines. There is no, however, universal definition for privacy, and many researchers have referred to the difficulties involved in trying to produce such a definition [37,38,39]. Newell, for example, argues that this is due to the multidisciplinary nature of privacy, and refers to the difficulty of relating studies from different disciplines [38]. In this section, we provide an overview of the span of definitions.

3.1. Privacy as a Right

Warren and Brandeis [40] 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” [41], 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. This basic definition remained the most famous as shown in the article titled “One hundred years of privacy” [42]. 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.

Solove [43] presented a taxonomy of privacy to serve as a framework for the development of the field of privacy law. In this taxonomy, he classified harmful privacy-related activities into four groups: *i*) Information collection: surveillance, interrogation. *ii*) Information processing: aggregation, identification, insecurity, secondary use, exclusion. *iii*) Information dissemination: breach of confidentiality, disclosure, exposure, increased accessibility, blackmail, appropriation, distortion. *iv*) Invasion: intrusion, decisional interference. He presented each of these activities and focused on how they affect an individual’s life in a harmful way, stressing on the importance of privacy as a right.

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.

3.2. Privacy as an Enabler

Improvements towards higher 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" [44] and Stefanos Gritzalis's 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"[45]. 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 [46].

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" [47], or simpler as "the ability to prevent other parties from learning one's current or past location" [48]. In [49] 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 [50], 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 [51] is an interesting tool to measure location privacy. Other types of privacy that need to be protected, as suggested in [23], include bodily privacy, territorial privacy, privacy of communications and information privacy.

Other approaches to the definition of privacy also include the idea of a free, uninfluenced decision making about oneself. Kupfer [52] 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".

Hong et al. [53] point out that privacy is a fluid and malleable notion with a range of needs and trust levels rather than being a single monolithic concept. They focused on empowering people with choice and informed consent, letting individuals share personal information with the right people and services, in the right situations, and at the right level of detail.

Altman [54] conceptualizes privacy as the "selective control of access to the self" regulated as dialectic and dynamic processes that include multimechanistic optimizing behaviors. Palen [55] argues that privacy management is not just about setting rules and enforcing them, rather than the continuous management between different spheres of action and degrees of disclosure within those spheres. DeCew [56] 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.

3.3. Privacy as a Commodity

Privacy seems to be a culturally relative right, but this doesn't mean that it is completely subjective [2]. For instance, privacy is more and more 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 to 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 agreement offers a convenient way of complying with the adequacy requirements of the EU Directive [57].

Privacy has nowadays become a commodity [58], 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 [59]. Consumers are becoming aware of the value of their personal data and are less and less ready to let businesses and companies, have this data without their explicit consent or even for free [60]. Trade-offs must therefore be established between, on the one hand, the consumers who want to obtain goods and services [61], and, on the other hand, businesses' eagerness to gather ever more knowledge and personal

data on consumers to better target and streamline their consumer base. Such tradeoffs are therefore becoming part of the requirements of AmI systems.

A very recent work by Li et al. [62] presents a theory of pricing private data. They provide a framework to estimate monetary value of private information queried with a certain degree of accuracy. They define as well the term “privacy budget” which refers to a limit on the quantity and/or accuracy of queries that any buyer can ask, in order to prevent an unacceptable disclosure of the data for the owner.

3.4. Discussion

Table 1 presents a categorization of the different approaches used to define privacy that have been discussed in this section. The table should be read as follows: We see privacy has been defined in three ways; as a right, as an enabler (to control personal data), and more recently as a commodity expressed through privacy policies of commercial products.

We note that this categorization is not strictly in a chronological order, in the sense that privacy was seen as a right in the 19th century, and still can be defined as such. Instead, Table 1 summarizes what can be seen as a chronological categorization of the value of privacy, as it has been understood or used. First introduced as an aspect of personal liberty, privacy then became an ability to control personal information, to being used today as a way to exchange personal information for a better service, or other commercial offers.

Figure 3 reflects the interest of each definition of privacy by showing the number of publications in the IEEE database that contain these keywords. Please refer to the annex for more information about the statistics. Privacy as a right is the oldest and most famous definition, then comes Privacy as an enabler. Finally, privacy as a commodity is a newest trending definition.

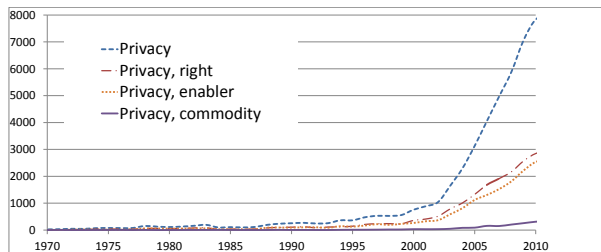


Fig. 3. Statistics about the number of papers containing the keyword privacy in the IEEE database over time.

The definitions of privacy have evolved. This has prompted the need to take privacy into consideration as a requirement for system design. Therefore, tradeoffs have to be found between privacy and other requirements, such as security, reliability and sociability.

For example, complete privacy would mean, roughly, sharing no personal information with anyone. Such behavior would exclude a person from social interactions with others and the environment. Typically, this is not the behavior that is expected from systems such as AmI and socio technical systems, where social interaction plays a crucial role.

Conversely, if a person was to share everything about her personal life with everyone, she may be considered very social, but risk her safety (physical, financial, etc.), as any mischievous person may use this information in harmful ways. Furthermore, in AAL, a privacy policy with extremely high constraints would not allow any personalisation of the system and therefore users would not benefit from its full potential.

4. PRIVACY ISSUES IN AAL

The pervasiveness of AAL technologies poses certain challenges to user privacy, due to the sometimes intrusive nature of the devices [64]. In this section, we discuss the privacy challenges which are emerging in AAL and illustrate them throughout with examples from our scenario. Moreover, we extend our discussion to the further challenges and issues typically encountered in AmI systems.

4.1. Identity Disclosure

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. In other cases they just carry an identification number.

Richter et al. described a way to remotely determine the nationality of an e-passport holder [65]. The detection is done at the logical level, by looking at the bytes that an e-passport sends as reply in response to some carefully chosen commands from the reader. The attack is able to distinguish e-passports from a wide set of nationalities. A potential abuse case that has been suggested, is an automated terrorist attack, for exam-

Right	Enabler	Commodity
<p>-to be let alone [40].</p> <p>-to be protected by law against the injury to the feelings, dignitary harm and reputational injury [43].</p> <p>-to keep a domain around us which includes all those things that are part of us, such as our body, home, property, thoughts, feelings, secrets and identity [63].</p>	<p>-to determine when, how, and to what extent personal information is communicated to others [44].</p> <p>-to control the ways in which personal information is obtained, processed, distributed, shared, and used [45].</p> <p>-to empower people with choice and informed consent to share personal information with the right people and services, in the right situations, and at the right level of detail [53]</p>	<p>-a non-monetary exchange of consumers' personal information for value such as higher quality service and personalized offers or discounts [59].</p> <p>-to empower individuals to control their private data through financial means [62].</p>

Table 1: Categorization of privacy definitions

ple a bomb explosion, designed to go off if someone holding a passport of a certain nationality comes close by.

4.2. Location Disclosure

4.2.1. Past geolocation Disclosure

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 [48], people often 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. RFID readers have been used in many AAL systems (e.g., [66]) to identify the location of elderly people even within a house and act accordingly.

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 ex-

ample, 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 whereabouts and habits, for marketing purposes, or even for surveillance reasons. [67]

4.2.2. Clandestine tracking

Clandestine tracking is a serious threat for Electronic passport holders as demonstrated by [67]; The standard for e-passport RFID chips (ISO 14443) stipulates the emission (without authentication) of a chip ID on protocol initiation. As this ID is different for every passport, it enables tracking the movements of the passport holder by unauthorised parties. Tracking is possible even if the data on the chip cannot be read. Clandestine tracking is a serious threat in AmI scenarios when the patient wears RFID tags or GPS devices.

Individuals carrying unique tags can be monitored and their location revealed if the monitoring agency knows the tags associated with those individuals [68]. A tag reader at a fixed location could track RFID-labeled clothes or banknotes carried by people passing by. Correlating data from multiple tag reader locations could track movement, social interactions, and financial transactions [69].

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 associated 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 that he has in his house or the one in Jane's house, then he could easily be identified at the place of this reader.

4.3. Inventorying

In his survey [70], Juels analyses how RFID raises a main privacy concern: inventorying. When an RFID tag has information about the manufacturer, or the cardholder, then the bearers of this tag are subject to inventorying. This means that an adversary could know, for example, 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.

4.4. Sensitive information disclosure

Lately, there has been a significant increase in the digital medical data being recorded by healthcare organisations. "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 organisations" [71].

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 [71]: *i.*) the data owner (the organisation 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 organisation 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-anonymised, they can be

combined with publicly available data and still identify the referred individual.

Another interesting source of privacy breach is 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 neighbours motivated by general curiosity, financial gain, child custody lawsuits and other considerations" [72].

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 authorisation to his HCS to take action. However, there could be a case, in which he does not push the button and the HCS realises 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, authorisation rules have to be predefined by Frank.

4.6. Profiling

Autonomic computing presumes autonomic profiling, which is defined in [73] as "a reiterative process of construction and application of profiles, entangling real time monitoring and real time M2M decision making".

Profiling problems can also be found in the energy consumption domain. Smart meters have unintended consequences for customer privacy [74]. Energy usage information stored at the meter and distributed thereafter acts as an information-rich side channel, exposing customer habits and behaviors. Certain activities, such as watching television, have detectable power consumption signatures. The customer has less control over the use of power information delivered to utility companies.

Profiling activity for the purposes of AAL is a continuous background activity; it includes extracting useful information (like user location, behaviour, room temperature), "enabling the identification of the user's needs, selecting suitable services and adjusting the pa-

rameters of the selected services in order to allow the AmI environment to behave according to the users' preferences, actions and expectations"[75].

Consequently, profiling activity is essential to meet user needs and preferences. In any case, such monitoring and surveillance may erode privacy and since large amounts of data may be stored, fears about personal data theft arise. The aim of AmI is to anticipate inferred habits and desires, which means that "one does not provide a profile on the basis of what one thinks to be one's preferences, but that one trusts the system to infer them and to adjust the environment accordingly" [73]. As discussed in the previous section (4.5), this seems to cause a loss of control, because the preferences and the resulting actions are not deliberately authorised by the final user.

In our scenario, profiling could lead to the same privacy breach as in section 4.5. Furthermore, Frank's personal data could also be compromised.

4.7. *Personal Data Matching*

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.

Entity resolution [76,77] is the process of identifying and merging references corresponding to the same real world entity. An *entity* could be a specific person, place, building, etc. Multiple references to a person, scattered through the Web, could be merged to form a single file, containing all the information that has been recorded about this person. This could also affect the privacy of the people that are somehow connected to this person, or even of groups of people.

Assume that two databases are linked, one containing medical records and the other containing social information. It could be deduced for example that the citizens of a specific area, or race have higher chances of having a contagious disease. If such knowledge is publicized, this could lead to stigmatisation or racist behavior against this group of people and diminish their chances of employment, or getting certain types of insurances [77]. Various real-world stories related to privacy and data matching have also been described by Clifton et al. [78] and by Frienberg [79].

It is discussed in [80] 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 [80] could identify the patient record of the governor of Massachusetts 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 [81], which presents a framework that analyses privacy and anonymity in social networks and re-identifies anonymised 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 become available to third parties, by data aggregation, without Frank's explicit authorization.

4.8. *Purpose Control*

Another aspect of privacy breach is discussed in [82]. Even if only authorised people access personal data, it is still not guaranteed that the personal data will be used for the intended purpose. Preventive mechanisms are not able to prevent a user to process data for other purposes after the same user has legitimately got access to them. [83] identifies the need of three additional elements, i.e. purpose, condition and obligation, besides the basic authorisation elements, i.e. subject, object and action.

Following our scenario, the local hospital gains access to Frank's Electronic Patient Record (EPR), in order to help him in his current emergency. However, there is no proof, or automated way (audit trail) to check that the hospital will not use Frank's EPR for other purposes, such as for research, or even selling it to a third party.

4.9. *Embarrassment and Social Isolation*

Ageing at home could have a positive impact on elderly population, especially when counter-balancing the negative aspects of institutionalisation (especially in the case of couples, who are used to living together and have done so autonomously and privately for decades. However, the use of assistive technology or leakage of disease data may cause embarrassment or even stigmatization [64]. Some AAL technologies evoked resistance in certain persons because

of their “handicapped-look” design. Moreover, an explicit feedback of information of a delicate nature coming from a device in front of other people or in public can be embarrassing for its user. Stigmatization can have major effects on the isolation of the user [64].

4.10. Divergent privacy requirements

An AAL environment typically involves multiple users with different and possibly divergent privacy requirements. In our scenario, the privacy preferences of Frank, Jane or any of their neighbours may be different; and it may not be possible to configure the system (e.g. the RFID readers) in such a way that all users are satisfied.

Addressing divergent privacy requirements is still an open problem in AmI, which becomes even more challenging in open environments such as public spaces, where the number and diversity of users is bigger, and each user may join or leave the environment at random times and without prior notice.

Very few recent studies have attempted to solve the problem either by satisfying the requirements of as many potential users as possible [84], or by resolving conflicts that arise from different privacy policies [85]. The proposed solutions, though, are still rather premature, and many gaps still remain mainly regarding the automatic adaptation of the system to the user and the contact, e.g. the automatic detection and resolution of conflicts.

4.11. Discussion

In this section we explore 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 perusal 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, such as his age, address, medical profile, marital status etc. Finally, Frank’s own HCS could take important decisions that would save

Frank’s life, without his specific approval at that very moment. Would that then be a privacy breach?

This example points out the moral and ethical issues raised by such systems. Indeed, in our scenario, Frank would have had to previously give his consent to such decision making process. This in turn would have been included in the system design, through so called “Informed consent” policies. A vast body of literature shows the importance of this area, particularly in the medical field with works such as [86,87,88,89].

To illustrate our discussion, Figure 4 shows the number of papers containing these privacy issues in the IEEE database. Please refer to the Annex for more details. We summarize our results by stating that the most studied issues, in fact accounting for 85 percent of papers, are: identity disclosure, sensitive information disclosure and location disclosure.

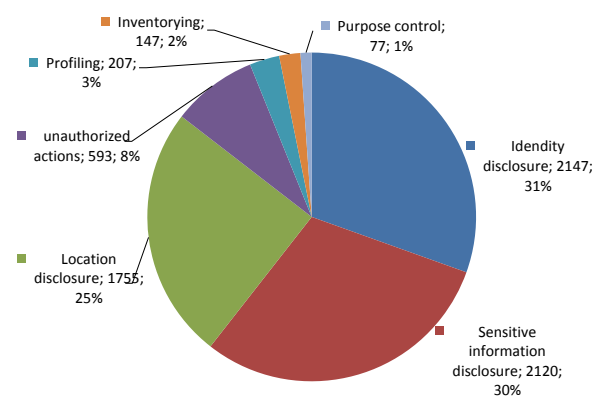


Fig. 4. Statistics about privacy issues in the IEEE database, number of publications and percentage.

5. PRESERVING PRIVACY

In this section we review the main approaches to preserve privacy. We first provide an overview of the solutions being proposed in fields with similar requirements such as the Web and then, discuss how they may be extended to address the additional privacy requirements of AAL.

Langheinrich provides different principles and guidelines for designing privacy-aware systems [18]. A summary of these guidelines is presented here:

1. Notice: When collecting data, it is crucial to notify users. For example, it is important to inform

- users that they are entering an area with video surveillance.
2. Choice and consent: In some situations, it is not enough anymore to simply announce and declare data collection - it is also required that collectors receive explicit consent from the data subject (by written contract, digital signature, check box.).
 3. Anonymity and pseudonymity: Collected data cannot be traced back to an individual, or at least unlinkable.
 4. Proximity and locality: In essence, information should not be disseminated indefinitely. For example, information collected in a building would stay within the building's network. Anybody interested in this information would need to be actually physically present in order to query it.
 5. Adequate security: Use of the best practice in securing communication, storage, etc.
 6. Access and Recourse: Trusting a system requires a set of regulations that separate acceptable from unacceptable behaviour with a reasonable mechanism for detecting violations and enforcing the penalties.

We have divided this section into two parts; the first one concerns data that can be accessed, but not in their original form. The second part mainly focuses on data access control; namely methodologies and techniques that enable users to control access to their personal data, e.g. which part of this data must be kept private, which may be shared and with whom.

There are also other approaches, such as auditing mechanisms [72], which could complement privacy policies. Additionally, various techniques to facilitate privacy-preserving data matching have been investigated, including one-way hash-encoding, secure multi-party computation approaches, such as commutative and homomorphic encryption and split data, Bloom filters, mapping attribute values into multi-dimensional spaces, and the use of public reference values. However, we believe that data matching is just one of the many privacy issues that arise in AmI, as shown in section 4, and hence do not further discuss these approaches in this paper. We strongly encourage the readers interested in this aspect of privacy, to read the excellent book of Christen [77].

5.1. Anonymising Personal Data

Privacy issues occur when someone's personal data becomes 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 neighbor, Frank, suffers from 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 from Alzheimer's disease. The property of being indistinguishable among a set of individuals is called anonymity. The problem is that even anonymised data can be combined and finally reveal 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.

Fung's et al. survey [90] provides a typical scenario for data collection and publishing (described in Figure 5). In the data collection phase, the data publisher collects data from record owners (e.g., Alice and Frank). In the data publishing phase, the data publisher releases the collected data to a data miner or to the public, called the data recipient, who will then conduct data mining on the published data. In this example, the hospital is the data publisher, patients are record owners, and the public is the data recipient. The data mining conducted at the medical center could be anything from a simple count of the number of men with Alzheimer to a sophisticated cluster analysis.




Fig. 5. Anonymization happens before the data publishing phase.

5.1.1. k -anonymity

A very popular approach during the last decade to anonymize personal data before releasing the collected data, has been based on the notion of k -anonymity [80]. In a k -anonymized dataset, each record is indistinguishable from at least $k - 1$ other records

with respect to certain identifying attributes [91]. k -anonymity can be achieved by suppressing and generalizing the attributes of users in the data [92]. Suppressing an attribute value means deleting it from the perturbed data and replacing it with a wildcard value that matches any possible attribute value.

Generalizing an attribute means replacing it with a less specific but semantically consistent value. One can see suppression as a special case of generalization, and that suppressing all attributes would guarantee k -anonymity. This is why a notion of utility in the data has to be incorporated whenever sanitizing data. Because the utility and privacy of data are intrinsically connected, no regulation can increase data privacy without also decreasing data utility [93]. The actual objective is to maximize utility by minimizing the amount of generalization and suppression [94]. Achieving k -anonymity by generalization with this objective as a constraint is a Non-deterministic Polynomial-time hard (NP-hard) problem [95].

5.1.2. ℓ -diversity

However, a more recent work [96], introducing ℓ -diversity, has proven that k -anonymity does not guarantee privacy against attackers using background knowledge, or when the sensitive data are lacking diversity.

Consider the example, in which a hospital publishes its 4-anonymized list of patients daily. Bob, suffering from cancer, is one of them, indistinguishable from at least three other patients in the daily list, since they all live in the same neighborhood and they all belong to the same age-group, 60-80. They are the only ones in the list who live in this neighborhood and by chance, they all have cancer. Alice, Bob's neighbor, saw an ambulance taking Bob to the hospital this morning. Alice can easily infer from the hospital's list that Bob has cancer, since she knows that he is 68 years old, he lives in this neighborhood and he was taken to the hospital that day, so he is one of the patients in that list. So, a list is ℓ -diverse if it contains at least ℓ "well-represented" values for the sensitive attribute S . Anatomy [97] is an example of a linear-time algorithm to compute tables that obey ℓ -diversity requirement.

5.1.3. t -closeness

But again, another more recent study [98], shows that ℓ -diversity may be difficult and unnecessary to achieve or insufficient to prevent attribute disclosure. For instance, suppose we are studying the test result for a particular virus (positive and negative). Further suppose that 99% of the results are negative, and only

1% being positive (have the virus). Then the two values have very different degrees of sensitivity. One would not mind being known to be tested negative, because then one is the same as 99% of the population, but one would not want to be known/considered to be tested positive. In this case, 2-diversity is unnecessary for an equivalence class that contains only records that are negative and on the other hand, might be difficult to achieve to cover someone being positive.

Another problem with the ℓ -diversity is that it does not take into account the semantical closeness of the sensitive values in each list. To tackle these issues, the notion of t -closeness was introduced in [98]. The idea behind it is to distribute the records of the table in lists where the distance between the distribution of a sensitive attribute in each list and the distribution of this attribute in the whole table are not higher than a threshold t . This way, the statistics of each anonymized-list regarding the sensitive attribute are "close" to the statistics of the full population. t -closeness is designed to reduce the information gain of an observer, which is the difference between the posterior belief (information acquired after seeing the released lists) and the prior belief (information before seeing the released lists).

5.1.4. Differential privacy

Dwork [99,100,101] introduces, and describes a mechanism achieving the notion of differential privacy. It proves that the formalization of Dalenius' [102] desideratum for statistical databases, that "nothing about an individual should be learnable from the database that cannot be learned without access to the database" cannot be achieved. Moreover, it is shown that even the privacy of someone not in the database can be at risk. For example, if we know that Frank has a twin sister, Mary, and there is a database containing Mary and her age, then we can infer Frank's age. Differential privacy, intuitively, is the additional privacy risk of someone's participation in a database.

5.2. Controlling access to personal data

Solutions to privacy preservation based on logic, mainly focus on permissions to access data (authorization problem). Following the definitions of privacy, it should be the people whose data is shared that decide who will have access to their personal data, either directly, e.g. by being asked, or by just being aware each time their data is broadcasted, or indirectly, by agreeing upon a privacy policy. In each case, they should al-

ways be in a position to control the flow of their personal data. Covington et al. [103] present a uniform access control framework that can be used to secure context-aware applications.

Unlike traditional systems where access control has been explored, access decisions may depend on the context in which requests are made. Moreover, context aware systems are gaining traction in literature as shown by the survey of Chen et al. [104] and the research of Al-Rabiaah et al. [105].

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 each context. For example, he could state that only his personal doctor can have access to his medical profile all the time and, in the case of an emergency, this access could also be granted to any other doctor in duty. This would prohibit Jane from viewing Frank's medical profile. We note that restricting access to personal data can also be enforced through security protocols, or devices (a simple lock could improve the protection of private files in a drawer). However, we find these approaches outside the scope of this work.

5.2.1. *Solutions specifically designed for AmI and AAL*

Although access control, as discussed above, has been identified as one of the most critical issues in Ambient Intelligence, few works so far have proposed specific methods and tools to address this problem. Many of the early works on privacy specification and access control in ubiquitous computing were based on the use of the P3P protocol. Langheinrich proposed a privacy-aware system that combines the P3P protocol with APPEL, an XML-based privacy specification language, and other privacy enhancing technologies such as privacy beacons and privacy proxies [106].

A similar approach, which combines P3P with security mechanisms for information hiding and encryption, was proposed in [107]. Soon, however, it was recognised that P3P, which was originally designed to support Web interactions typically involving e-commerce and business applications, could not meet the requirements of most context-aware applications. Studies such as [108,109] proposed extending the P3P protocol and the APPEL language with appropriate tags and datatypes that better fit the specific characteristics of Ambient Intelligence systems.

Most of the recent works in the field are based on the use of rule languages for the specification of access control policies. In [110], for example, a semantically

rich, policy-based framework that constrains the information flow in a context-aware system is built on top of a Web Ontology Language (OWL)⁵ ontology that represents dynamic aspects of context-aware systems and a combination of OWL-DL (a sub-language of OWL) and Jena⁶ rules specifying the policy to perform reasoning.

The framework enforces user's privacy preferences using static information about the user as well as dynamic information observed and inferred from the context. In this case, privacy preferences are actually "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.

Similar approaches were also adopted by earlier works in the same field: a combination of an OWL ontology and Jess⁷ rules for the specification of privacy policies were proposed in [111]; a logic-based approach based on First-Order Logic proposed in [112]; and a similar rule-based approach using Event-Condition-Action rules presented in [113].

Being based on logic and rule languages, such approaches combine several desirable features, such as simplicity and flexibility (rule languages are easy to write and understand as the use a natural-language-like syntax), formality (they have well-defined syntax and semantics), reasoning support and high-level abstraction. Their main limitations are two: (a) being based on classical logics, they cannot deal with uncertainty, missing data and ambiguities, which are typical features of Ambient Intelligence environments; and (b) they are all based on the assumption of a centralised or semi-centralised architecture (such as the server-client model), which is not adopted by most Ambient Intelligence systems.

A first infrastructure to facilitate the development of privacy-sensitive ubiquitous computing applications is the Context Fabric (Confab) [53]. It is tailored for context-aware computing [114], a common aspect of ubiquitous computing in which sensors and other data sources are leveraged to provide computing systems with an increased awareness of a user's physical and

⁵<http://www.w3.org/TR/owl-features/>

⁶<http://jena.apache.org/index.html>, a Java framework for building Semantic Web and Linked Data applications

⁷<http://herzberg.ca.sandia.gov/>, a rule engine for the Java platform

social environment. Confab offers users the ability to select different privacy settings when sharing specific data. For instance, the users have an option to specify how long the data they share should be retained before being deleted; and another option to specify the maximum number of previous values that can be retained.

An approach that is closer to the real needs and requirements of Ambient Intelligence environment was the DEAL language, proposed in [5]. DEAL is a formal high-level authorisation language, aiming to specify access control policies in open and dynamic distributed systems. It was designed as an access-control and authorisation layer on top of Contextual Defeasible Logic, which is a distributed representation model specifically designed for AmI environments [115]. Being based on Defeasible Logic, it supports reasoning with missing or ambiguous information and conflicting privacy policies. It also uses two types of preferences - a rule priority relation and a preference order on information sources - to resolve conflicts between conflicting policies. Its main limitation is the assumption that context knowledge must be encoded in Defeasible Logic, which makes its use as a general-purpose solution for AmI and AAL environments rather problematic.

5.2.2. General-purpose solutions

Instead of developing data access control solutions for Ambient Intelligence systems from scratch, an alternative approach would be to reuse methodologies, tools and guidelines that have been designed for other domains with similar requirements and needs, such as the Web, and adapting them to the special characteristics of Ambient Intelligence environments. Below, we review some general-purpose solutions, most of which are based on logic-based methodologies for the specification of access control policies.

PROTUNE (Provisional Trust Negotiation) [116] is a system for specifying and cooperatively enforcing security and privacy policies. It 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 provides the advantage of well-defined semantics and machine interoperability, hence allowing for automated negotiations. Furthermore, it enables a straightforward integration with the ontology-based models for context representation that are used by most current AmI and AAL systems.

In PROTUNE, policies are sets of Horn rules, on which the system has to perform several kinds of sym-

bolic 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 also 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. This is a key requirement for the establishment of trust in the interactions between users and Ambient Intelligence systems, which is essential for the adoption and success of such systems. Regarding its deployment in AmI and AAL environments, the main limitation of PROTUNE is its inability to deal with the several types of uncertainty that characterise such environments, e.g. incomplete, imprecise or ambiguous context data and unreliable wireless connections between devices.

Casper [117] is a framework in which mobile users can entertain location-based services without the need to disclose their exact private location information. Mobile users register with Casper by a user-specified privacy profile. Casper has two main components, the location anonymiser and the privacy-aware query processor. The location anonymiser acts as a third trusted party that blurs the exact location information of each user into a cloaked spatial area that matches the user privacy profile. The privacy-aware query processor is embedded into traditional location-based database servers to tune their functionalities to be privacy-aware by dealing with cloaked spatial areas rather than exact point information. Another approach to protect location privacy, while providing exact location-based services, is by frequently changing pseudonyms of users [48].

The logic-based approaches proposed in [118,119, 120] enable the representation and reasoning about policies, credentials, and requests in distributed authorisation. In all these works, authorisation is defined as "the process of specifying an access control policy that is used to determine whether a requester, with a given valid identity, is permitted to consume a particular requested service". They also support delegation of authorisation, enabling entities to delegate the authority over an attribute to another entity, i.e., the entity trusts another entity's judgment about the attribute.

This feature is particularly important for domains such as the Web (that they were originally designed for) but also for Ambient Intelligence, where many of the interacting entities are unknown to each other, and

often there is no central authority that everyone trusts. DEAL, the distributed authorisation language for Ambient Intelligence [5] that we discussed in the previous subsection, was based on the same ideas, adapting them to the special features of AmI and AAL environments.

The use of modal logics, such as deontic logic, has also been proposed by recent studies for the specification of privacy policies, but also for verifying privacy policies with respect to high-level regulations applied to the whole system. Aucher et al. [22,121] 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. For example, the privacy policy not *Permitting* an agent a to *Know* at the same time both pieces of information p and q may be expressed as follows: $\neg PK_a(p \wedge q)$. The logic of [121] then adds dynamics to this statement, by using the $[send \phi]$ operator, which reads as “information ϕ is sent to the agent”. This addition allows the inference of the formula $K_a p \rightarrow \neg P[send q]$, read as: “if agent a Knows p , then sending q to agent a is not Permitted”.

The Coprelabri (*computers and privacy regulations: the logical bridge*) project [22], is built on top of a logical language that can be used to represent and reason on privacy policies. It may be used to provide writing assistance to lawyers in charge of specifying privacy policies, regulations and law. It may also be used to check that a given policy is compliant with a set of high-level regulations. Deontic Logic for Privacy (DLP logic) [122] 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, which are both very relevant to Ambient Intelligence. Kanovich et al. [123] 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 assumed 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.

Among the several non-logical approaches that have been proposed for the problems of privacy preservation and data access control, we distinguish trust management systems, such as PolicyMaker [124], KeyNote [125], REFEREE [126] and SPKI/SDSI [127], which focus on aspects of trust, such as access control and authorization; and the *Personal Data Stream (PDS)* framework proposed in [128]. PDS is designed to give users new data management tools, based on three foundational design principles: privacy 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. Compared to logic-based solutions, the main limitations of such approaches are the lack of formal declarative semantics and the limited expressive capabilities.

5.3. Discussion

We first want to emphasise again that the two main approaches presented in this section to preserve privacy, serve the different purposes. Namely, data masking / anonymising is typically used when personal data is 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’ privacies, would use one of these solutions. In other words data masking is all about **what** kind of data third parties have access to.

On the other hand, the data access control problem is about deciding **who** has access to personal, typically not anonymised data. Although, it has been identified as one of the most critical issues in Ambient Intelligence, and particularly AAL, there is no general framework yet; only domain-specific solutions designed for the specific needs and requirements of individual systems. The logic-based approaches proposed in [110,112,111,113], for example, are based on simplifying assumptions, such as perfect knowledge of

context, which are not valid in Ambient Intelligence environments. DEAL [5] may not serve as a generic solution as well, as it makes specific restricting assumptions about the representation model.

The general-purpose solutions that we presented in 5.2.2 address different facets of the same problem. Frameworks, such as Protune [116], provide ways to semantically represent and reason about privacy policies, enabling also the provision of explanations as a means for the establishment of trust, which is a key requirement for the success of Ambient Intelligence systems such as AAL systems. Approaches, such as those proposed in [118,119,120], enable distributed authorisation and delegation, which are also essential given the open, dynamic and distributed nature of AmI environments. The modal-logics-based approaches of [22,121,122] enable a form of centralised control by verifying distributed privacy policies with respect to a set of norms applied to the overall system. Combining all these different aspects in a single logic-based privacy framework will lead to a general-purpose framework for Ambient Intelligence and AAL.

A combination of anonymising and authorisation approaches would also be useful, as it would enable users to have control of any type of data, including anonymised and encrypted data, adding another layer of privacy control. Figure 6 shows the number of papers studying these privacy preservation techniques in the IEEE database; refer to the Annex for more details. k -anonymity is the most popular.

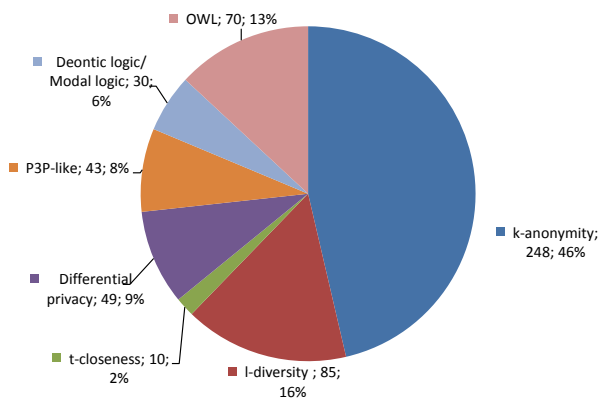


Fig. 6. Statistics about privacy preserving techniques in the IEEE database, number of publications and percentage.

6. LESSONS LEARNED, GAPS, AND FUTURE PERSPECTIVES IN THE AAL DOMAIN

In this section, we draw the lessons learned from the fields of AAL, highlight the gaps we identified, and present insights on the future of privacy in AmI systems and AAL.

6.1. Defining and quantifying privacy

First, we observe that, even though privacy has been a concern for people for over a century, it is a concept that has not yet been formally defined in a way that would be globally accepted. Moreover, depending on the chosen definition, privacy can be breached in multiple manners. (i.e. location privacy, patient privacy, profiling etc). In the evolution of the concept, privacy is becoming a commodity, exchanged for better, more personalised services. Privacy policies, preferences and agreements replace data protection legislations.

Then, we note that, except for the recent work by Alessandro Acquisti et al. [129], there is little research done on how to generally quantify privacy and more specifically how to measure the monetary value of private information, for example when it is somehow anonymised. An interesting open question is regarding the trade-offs between on one hand utility and privacy, and on the other hand privacy and risk. Furthermore, there is a lack of tools and frameworks to allow users to trade, in a public market, their private information collected in AmI and AAL systems, for some goods and services with third parties.

Furthermore, most of the current work in the privacy domain focuses on how to anonymise the collected data before publishing it [93]. Anonymisation however, is not enough to protect vital personal information. Anonymised data can easily be aggregated and combined (e.g. by data mining techniques or simply by knowing some personal information about someone) and thus become de-anonymised. Furthermore, authorisation does not guarantee that if data is leaked, it will not be harmful to the individuals. Even if anonymisation and authorisation policies work properly, it should be guaranteed that any personal information is used for the intended purpose by those who have access to it.

6.2. Principles and guidelines for privacy preservation

We recall that a decade ago, Langheinrich tried to build a roadmap for AmI system designers, who want

to respect their users' privacies, and updates need to be made to fill the gaps. Anonymity, as well as security, have been extensively studied. However, only a small percentage of the literature deals with the issues of profiling, inventorying, and purpose control, linked to Lanheinrich's sixth principle regarding access and recourse. More specifically, this sixth principle should encourage AmI and AAL systems designers to:

- only collect data for a well-defined purpose (no in-advance storage),
- only collect data relevant for the purpose (not more), and
- only keep data as long as it is necessary for the purpose.

However, and as mentioned above, privacy is increasingly seen as a commodity, making the observance of the access and recourse principle unfeasible. Using private data for more than one purposes, even if this is not stated in the initial user's agreement, as well as keeping the private data for as long as possible, are all principles for maximising the profit of someone who wants to monetise privacy. Hence, the principle of access and recourse is in conflict with the perspective of privacy as a commodity.

In order to preserve and respect this principle, privacy agreements should become easier for users to read and understand, and more flexible for them to use. For example, instead of asking users whether they accept all the terms of the policy agreement, users should also have the ability and the means to select some of them, or even to include additional ones. This would be particularly useful for AAL systems, as users typically have different kinds of medical conditions, which need specific attentions and requirements.

This brings up the importance of creating guidelines for AmI in general and AAL systems in particular, now. They would include the development of "privacy friendly" operating systems and devices, so that smart devices do not remain "black boxes" to their owners. In fact, no one wants these now ubiquitous devices that we carry with us everywhere, from our home to our offices, to be their ultimate spies. Examples of such spying applications are numerous, e.g. the ones that take users' past locations to determine whether some users have spent time in the same place. For AAL users, the dangers are even greater, as they do not always realise the risks they run when they allow for their personnel information to be collected. Often not aware enough of the underlying technological meanders, such users

often do not understand all the commitment layers of the services they are being proposed.

Moreover, with the increase of personal and corporate data stored in mobile devices, sharing such data requires preventive security and privacy mechanisms to protect and regulate access to this data shared by other users. Besides protecting the private data on the device, the protection of data used by the services the device connects to must be considered and enforced. Protection against insider as well as against outsider threats can for example be ensured with authentication to services and privacy-friendly biometrics with protected biometrics templates.

6.3. *Tools and frameworks for privacy protection*

In fact, when using third-party services, a large quantity of personal information is shared. Therefore, tools to analyse the amount of personal data revealed and to validate the privacy-protecting properties of communication protocols must be used. Collaborative and shared applications must ensure privacy and confidentiality, for example when using calendar events, and should be independent of any third-party or big company servers. In the AAL field, a number of applications show the advantages of connecting the users' calendars to a number of other events such as outpatients medical visits, friends and family smart phone applications and smart home applications for appliances. However, the complexity of such systems and complications due to the lack of standards is too often minimised.

Another gap we identified is the lack of general frameworks and tools to empower users to select the level of anonymisation that they would like to have for their private information before publishing their data. Indeed, currently, users have little or no choice for selecting privacy levels, and it is up to the data holders, such as hospitals, web service providers and regulators to select these levels. In the rare case where a framework exists, it is usually limited to a specific application domain (usually health care), or about a specific data entity (GPS location [117]). It is worth noting that many potential users of AAL systems are not health care patients. They are people who are healthy enough to live alone but, because of their advanced age, they have reduced capabilities.

AmI is very special regarding privacy issues, and this is mainly due to profiling. Indeed, users of AmI and AAL systems want these systems to act according to their expectations and needs. In order to do this,

AmI systems need as much information as they can get from users. Hence, they collect, store, and may even distribute this information, according to certain situations and needs, and this is exactly where the danger for users' privacy lays.

A crucial aspect of the problem of privacy in AAL is therefore to make privacy usable for users; it should be user-friendly, particularly considering the fragility of its users, and allow users to cooperate with each other and with the system [130]. There must be a trade-off between how much privacy the system allows, and how cooperative and convivial the system can be [131]. On the one hand, the more private the AAL system is, the less opportunities it offers to its users to cooperate and exchange information with each other, and the less convivial it is for users [132,133]. On the other hand, AmI and AAL systems that share all their information with all other users offer better cooperation, but minimise their users' privacy.

6.4. Privacy-aware system design

Finally, although the awareness for privacy is increasing, it is still a highly underestimated aspect in the design phase of AmI and AAL systems. Paradoxically, it is one of the main concerns of end users. Hence, for AmI systems in general to gain some ground and be part of our lives, as they aim to become, privacy protection should be included in the first steps of the design.

Privacy By Design approaches provide many positive examples of privacy protection, with "case studies to show that systems with no personal data – or at least with much less personal data – could have the same functionalities" [134], or that incorporating user-centred design practices enable users to make wise choices [135]. Moreover, privacy policies and contracts must be transparent, readable and user-configurable. This will help users trust such systems and focus on what they can bring to them. For AAL users, these functionalities are even more important due to the fact that they often are not very technically literate.

Of course, it is not sufficient that system designers convey their intentions via the use of privacy policies. Privacy should also be ensured through strong security protocols. Similarly, a key-question is the accountability/awareness of the user to set up the level of privacy that suits her best, while still keeping the system working as intended.

Clearly, there is still a gap in the literature, both from legal and technical perspectives, to determine how to adapt this level of privacy and decide who has the right/duty to deliver the information. There is a need for research to, on the one hand, focus on increasing the users' awareness and accountability by displaying the risk(s) related to some information disclosure or the value of this information. Of course this shouldn't be at the expense of the simplicity of use and user-friendliness of the system, service or application.

On the other hand, there is also a need for research related to determining who has the right to relax or disclose specific information with respect to someone's mental and physical state, e.g., age, dementia. So, while in principle the system must be user-centric, in practice this principle has to be adapted case by case, and with appropriate regulations. Particularly relevant is the domain of application and the target users, for example in the case of AAL.

7. CONCLUSION

With the growing number of the ageing population and the pervasiveness of Ambient Intelligence (AmI) and Ambient Assisted Living (AAL) systems, privacy issues have become key challenges. In this article, we raise the question of which specific issues and solutions currently exist in AmI and AAL environments, and how are future perspectives shaping up.

First, to understand how, up to now, the concept of privacy has evolved, 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" [41,40], to enabling "control over personal information" [52,56,44], to recently becoming a powerful "commodity" to be traded [59].

Second, building on our previous work, in which we introduced this new tripartite categorisation of privacy as a right, an enabler, and a commodity, we highlight the specific privacy issues increasingly being raised in AAL. To clarify this evolution, we extend a categorisation of privacy definitions following these distinctions, and present a graph illustrating the evolution of the concept of privacy over the last forty years. The statistics we present are based on the number of papers, part of the IEEE database of publications, containing as main keyword the word *privacy*. It is worth noting, over the last ten years, the phenomenal increase in the use of the term.

Section	Lessons learned	Perspectives
Privacy definition	Previously, privacy was mainly considered as a right Now, privacy tends to be considered as a commodity	New development of tools, measures, framework for privacy as a commodity in AmI and AAL
Privacy issues	Privacy in AmI-AAL mature issues: identity disclosure, sensitive information disclosure and location disclosure	Explore purpose control, profiling, and inventorying issues, which are not yet mature
Privacy preserving	Privacy in AmI-AAL mature issues: anonymisation techniques and privacy policies Langheirich design principles need updates to add the aspect of privacy as a commodity and allow users to release their private data in exchange of money or services	Development of privacy in AmI frameworks combining both anonymisation techniques and privacy policy. Development of trade-offs between privacy, user-friendliness of AmI and AAL systems.

Table 2: Overview of lessons learned and future perspectives.

Third, we review and discuss the most common privacy issues pertaining to AmI and AAL environments. We find that geolocation may give rise to very serious privacy breaches, especially when location logs are kept. Furthermore, while RFID promises interesting identification techniques, if RFID security is not carefully designed, the information stored in RFID tags may seriously be at risk. Indeed, using sophisticated data aggregation methods, AAL patient data can be acquired by third parties, such as health care providers and telecommunication services, even from publicly available anonymised data.

Fourth, personal data, like any other data, can be leaked and become publicly available, while unauthorised actions and decisions by a third person represent yet other types of privacy breaches. Profiling is singled out as one of the key aspects specific to AmI systems, making it easier to target someone or a whole category of people. This brings up another dimension of privacy, one which is concerned with the actual use of the data by the people who have been granted access to it. It is worth noting that among the most mature areas are: identity disclosure, sensitive information disclosure and location disclosure. Furthermore, it is foreseen that future perspectives in the areas considered not mature yet, include: the exploration of purpose control, profiling, inventorying issues, and divergent privacy preferences. Due to their often minimal technical knowledge, the AAL users represent a particularly exposed segment of the population.

Fifth, we describe the two most common approaches used to preserve privacy, for example by modifying the available data and providing authorisation mechanisms. Such approaches address two different questions, namely **what** kind of data can a third party acquire and **who** can acquire private data, respectively.

We observe that the most advanced areas lay in anonymisation techniques and privacy policy. Furthermore, we note that Langheirich's design principles need to be updated to adapt to the more recent aspect of privacy as a commodity and allow users to release their private data in exchange of financial gain or services. We also foresee future perspectives in the development of, on the one hand, frameworks combining both anonymisation techniques and privacy policy, and, on the other hand, trade-offs between privacy and user-friendliness in AmI and AAL systems. The evolution of smart homes, the increasing number of people needing recurrent help at home and the difficulty of the institutional healthcare systems to cope with this situation are all contributing to the importance of research and developments in this domain.

Finally, we provide some insights on the lessons learned, particularly the challenges brought by profiling and techniques such as anonymisation. We discuss alternative approaches to addressing privacy issues, namely by establishing tradeoffs between concerns of privacy and concerns of conviviality and usability in order to provide better cooperation among users, particularly with those with low level of technical literacy typically part of the senior population, and between

users and the system. We highlight the gaps we have identified from our literature reviews and through our statistical analysis of the publications included in the IEEE database, we then contrast these gaps with the more mature issues. Throughout our article, we illustrate how these issues may arise and these techniques be used in real-lives situations, particularly those set in the AAL domain, with a motivating scenario validated by the HotCity of Luxembourg. We also believe that privacy in AmI and specifically in AAL systems needs to be addressed in order to avoid potential harm to an increasing population whose members expect that such systems will work for their users, as they should, rather than against them. Hence, great opportunities lay ahead, but they have to be grasped, and in timely manner.

References

- [1] C. Edge, W. 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.
- [2] A. D. Moore, Privacy: Its meaning and value, *American Philosophical Quarterly* 40 (2003) pp. 215–227.
- [3] D. L. Métayer, Privacy by design : towards a systematic approach, *Atelier Protection de la Vie Privée (APVP 2011)*, 2011.
- [4] D. J. Cook, J. C. Augusto, V. R. Jakkula, Ambient intelligence: Technologies, applications, and opportunities., *Pervasive and Mobile Computing* (2009) 277–298.
- [5] I. Genitsaridi, A. Bikakis, G. Antoniou, Deal: A distributed authorization language for ambient intelligence, *IJACI* 3 (2011) 9–24.
- [6] J. García-Vázquez, Design of ambient information systems to assist elderly with their medication compliance, *Journal of Ambient Intelligence and Smart Environments* 5 (2013) 657–659.
- [7] G. Bleser, D. Steffen, M. Weber, G. Hendeby, D. Stricker, L. Fradet, F. Marin, N. Ville, F. Carré, A personalized exercise trainer for the elderly, *Journal of Ambient Intelligence and Smart Environments* 5 (2013) 547–562.
- [8] A. Bono-Nuez, R. B. Marín, R. Casas, B. Martín-del-Brío, Ambient intelligence for quality of life assessment, *Journal of Ambient Intelligence and Smart Environments* 6 (2014) 57–70.
- [9] J. Doyle, A. Kealy, J. Loane, L. Walsh, B. O’Mullane, C. Flynn, A. MacFarlane, B. Bortz, R. B. Knapp, R. Bond, An integrated home-based self-management system to support the wellbeing of older adults, *Journal of Ambient Intelligence and Smart Environments* 6 (2014) 359–383.
- [10] R. Igual, I. Plaza, C. Medrano, M. A. Rubio, Personalizable smartphone-based system adapted to assist dependent people, *Journal of Ambient Intelligence and Smart Environments* 6 (2014) 569–593.
- [11] W. Sanchez, A. M. Rebollar, W. Campos, H. Estrada, V. Pelechano, Inferring loneliness levels in older adults from smartphones, *Journal of Ambient Intelligence and Smart Environments* 7 (2015) 85–98.
- [12] R. Velik, A brain-inspired multimodal data mining approach for human activity recognition in elderly homes, *Journal of Ambient Intelligence and Smart Environments* 6 (2014) 447–468.
- [13] N. P. Pulido, J. A. López-Riquelme, J. F. Melero, M. Á. V. Rodríguez, A. J. Barrios-León, A service robot for monitoring elderly people in the context of ambient assisted living, *Journal of Ambient Intelligence and Smart Environments* 6 (2014) 595–621.
- [14] G. Moretti, S. Marsland, D. Basu, G. S. Gupta, Towards a monitoring smart home for the elderly: One experience in retrofitting a sensor network into an existing home, *Journal of Ambient Intelligence and Smart Environments* 5 (2013) 639–656.
- [15] P. Brey, Freedom and privacy in ambient intelligence, *Ethics and Information Technology* 7 (2005) 157–166.
- [16] J. Bohn, V. Coroama, M. Langheinrich, F. Mattern, M. Rohs, Social, economic, and ethical implications of ambient intelligence and ubiquitous computing, *Ambient intelligence* (2005) 5–29.
- [17] A. Cavoukian, Privacy by Design: Leadership, Methods, and Results., in: *European Data Protection*, Springer, 2013, pp. 175–202.
- [18] M. Langheinrich, Privacy by design - principles of privacy-aware ubiquitous systems, in: G. Abowd, B. Brumitt, S. Shafer (Eds.), *Ubicomp 2001: Ubiquitous Computing*, volume 2201 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, 2001, pp. 273–291. URL: http://dx.doi.org/10.1007/3-540-45427-6_23.
- [19] Y. Wang, A. Kobsa, Privacy-enhancing technologies, in: *Handbook of Research on Social and Organizational Liabilities in Information Security*, Information Science Reference, 2008, pp. 352–375.
- [20] C. A. Brodie, C.-M. Karat, J. Karat, An empirical study of natural language parsing of privacy policy rules using the sparcle policy workbench, in: *Proceedings of the second symposium on Usable privacy and security, SOUPS ’06*, ACM, New York, NY, USA, 2006, pp. 8–19. URL: <http://doi.acm.org/10.1145/1143120.1143123>. doi:10.1145/1143120.1143123.
- [21] Organization for the Advancement of Structured Information Standards., Privacy policy profile of XACML v2.0, Technical Report, 2005.
- [22] G. Aucher, C. Barreau-Saliou, G. Boella, A. Blandin-Obernesser, S. Gams, G. Piolle, L. Van Der Torre (Coprelabri), The Coprelabri project : the logical approach to privacy, in: *2e Atelier Protection de la Vie Privée (APVP 2011)*, Sorèze, France, 2011. URL: <http://hal-supelec.archives-ouvertes.fr/hal-00606014>.
- [23] M. Karyda, S. Gritzalis, J. H. Park, S. Kokolakis, Privacy and fair information practices in ubiquitous environments: Research challenges and future directions, *Internet Research* 19 (2009) 194–208.
- [24] S. Lanspery, J. Callahan Jr, J. Miller, J. Hyde, Introduction: staying put, *Staying put: Adapting the places instead of the people* (1997) 1–22.

- [25] L. E. Hebert, P. A. Scherr, J. L. Bienias, D. A. Bennett, D. A. Evans, Alzheimer disease in the us population: prevalence estimates using the 2000 census, *Archives of neurology* 60 (2003) 1119.
- [26] R. L. Ernst, J. W. Hay, The us economic and social costs of alzheimer's disease revisited., *American Journal of Public Health* 84 (1994) 1261–1264.
- [27] P. Grayons, Technology and home adaptation, *Staying Put: Adapting the Places Instead of the People* (1997) 55–74.
- [28] P. Caire, L. van der Torre, Convivial ambient technologies: Requirements, ontology and design, *Comput. J.* 53 (2010) 1229–1256.
- [29] V. Efthymiou, P. Caire, Privacy challenges in ambient intelligent systems: A critical discussion, in: *Atelier Protection de la Vie Privee*, Island of Groix, February 29, 2012, 2012.
- [30] H. Jonker, J. Pang, Bulletin boards in voting systems: Modelling and measuring privacy, in: *ARES, IEEE*, 2011, pp. 294–300.
- [31] J. Benaloh, M. Chase, E. Horvitz, K. Lauter, Patient controlled encryption: ensuring privacy of electronic medical records, in: *Proceedings of the 2009 ACM workshop on Cloud computing security, CCSW '09*, ACM, New York, NY, USA, 2009, pp. 103–114. URL: <http://doi.acm.org/10.1145/1655008.1655024>. doi:10.1145/1655008.1655024.
- [32] A. Juels, R. L. Rivest, M. 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*, ACM, New York, NY, USA, 2003, pp. 103–111. URL: <http://doi.acm.org/10.1145/948109.948126>. doi:10.1145/948109.948126.
- [33] N. Li, N. Zhang, S. K. Das, B. Thuraisingham, Privacy preservation in wireless sensor networks: A state-of-the-art survey, *Ad Hoc Networks* 7 (2009) 1501 – 1514. Privacy and Security in Wireless Sensor and Ad Hoc Networks.
- [34] S. Jones, S. Hara, J. Augusto, eFRIEND: an ethical framework for intelligent environments development, *Ethics and Information Technology* 17 (2015) 11–25.
- [35] V. Efthymiou, P. Caire, Diagram Analysis Report: Use Cases for Conviviality and Privacy in Ambient Intelligent Systems, University of Luxembourg, SnT, Luxembourg, 2012.
- [36] P. G. Frankl, E. J. Weyuker, Testing software to detect and reduce risk, *Journal of Systems and Software* 53 (2000) 275–286.
- [37] J. Burgoon, Privacy and communication, *Communication Yearbook* 6 (1982) 206–249.
- [38] P. Brierley-Newell, A cross-cultural comparison of privacy definitions and functions: A systems approach, *Journal of Environmental Psychology* 18 (1998) 357–371.
- [39] H. Leino-Kilpi, M. Välimäki, T. Dassen, M. Gasull, C. Lemonidou, A. Scott, M. Arndt, Privacy: a review of the literature, *International Journal of Nursing Studies* 38 (2001) 663–671.
- [40] S. D. Warren, L. D. Brandeis, The right to privacy, *Harvard Law Review* 4 (1890) pp. 193–220.
- [41] T. Cooley, *A Treatise on the Law of Torts: Or the Wrongs which Arise Independent of Contract*, Callaghan, 1878. URL: <http://books.google.lu/books?id=UgsrAQAAMAAJ>.
- [42] K. Gormley, One hundred years of privacy, *Wis. L. Rev.* (1992) 1335.
- [43] D. J. Solove, A Taxonomy of Privacy, *University of Pennsylvania Law Review* 154 (2006) p477+.
- [44] A. Westin, *Privacy and Freedom*, New York Atheneum, 1967.
- [45] S. Gritzalis, Enhancing web privacy and anonymity in the digital era, *Inf. Manag. Comput. Security* 12 (2004) 255–287.
- [46] D. Wright, S. Gutwirth, M. Friedewald, P. De Hert, M. Langheinrich, A. Moscibroda, Privacy, trust and policy-making: challenges and responses, *Computer Law & Security Report* 25 (2009).
- [47] M. Duckham, L. Kulik, Location privacy and location-aware computing, *Dynamic & mobile GIS: investigating change in space and time* (2006) 34–51.
- [48] A. R. Beresford, F. Stajano, Location privacy in pervasive computing, *IEEE Pervasive Computing* 2 (2003) 46–55.
- [49] M. P. Scipioni, M. Langheinrich, Towards a new privacy-aware location sharing platform, *Journal of Internet Services and Information Security* 1 (2011).
- [50] J. Krumm, A survey of computational location privacy, *Personal Ubiquitous Comput.* 13 (2009) 391–399.
- [51] R. Shokri, G. Theodorakopoulos, J.-Y. Le Boudec, J.-P. Hubaux, Quantifying location privacy, in: *Proceedings of the 2011 IEEE Symposium on Security and Privacy, SP '11*, IEEE Computer Society, Washington, DC, USA, 2011, pp. 247–262. URL: <http://dx.doi.org/10.1109/SP.2011.18>. doi:10.1109/SP.2011.18.
- [52] J. Kupfer, Privacy, autonomy, and self-concept, *American Philosophical Quarterly* 24 (1987) pp. 81–89.
- [53] J. I. Hong, An architecture for privacy-sensitive ubiquitous computing, in: *In MobiSYS 2004: Proceedings of the 2nd international conference on mobile systems, applications, and services*, ACM Press, 2004, pp. 177–189.
- [54] I. Altman, Privacy Regulation: Culturally Universal or Culturally Specific?, *Journal of Social Issues* 33 (1977) 66–84.
- [55] L. Palen, P. Dourish, Unpacking "privacy" for a networked world, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '03*, ACM, New York, NY, USA, 2003, pp. 129–136. URL: <http://doi.acm.org/10.1145/642611.642635>. doi:10.1145/642611.642635.
- [56] J. DeCew, *In Pursuit of Privacy: Law, Ethics, and the Rise of Technology*, G - Reference, Information and Interdisciplinary Subjects Series, Cornell University Press, 1997. URL: <http://www.google.lu/books?id=GjfkK56dtlAC>.
- [57] G. Steinke, Data privacy approaches from us and eu perspectives, *Telematics and Informatics* 19 (2002) 193 – 200. Regulating the Internet: EU and US perspectives.
- [58] S. G. Davies, Re-engineering the right to privacy: how privacy has been transformed from a right to a commodity, in: *Technology and privacy*, MIT Press, 1997, pp. 143–165.
- [59] M. J. Culnan, R. J. Bies, Consumer privacy: Balancing economic and justice considerations, *Journal of Social Issues* 59 (2003) 323–342.
- [60] C. J. Dommeyer, B. L. Gross, What consumers know and what they do: An investigation of consumer knowledge, awareness, and use of privacy protection strategies, *Journal of Interactive Marketing* 17 (2003) 34–51.
- [61] Z. Papacharissi, Privacy as a luxury commodity, *First Monday* 15 (2010) 2–5.
- [62] C. Li, D. Y. Li, G. Miklau, D. Suci, A theory of pricing private data, in: *Proceedings of the 16th Interna-*

- tional Conference on Database Theory, ICDT '13, ACM, New York, NY, USA, 2013, pp. 33–44. URL: <http://doi.acm.org/10.1145/2448496.2448502>. doi:10.1145/2448496.2448502.
- [63] Y. Onn, M. Geva, et al., Privacy in the digital environment, Haifa Center of Law & Technology (2005) 1–12.
- [64] P. Novitzky, A. F. Smeaton, C. Chen, K. Irving, T. Jacquemard, F. O Brolchain, D. O Mathuna, B. Gordijn, A review of contemporary work on the ethics of ambient assisted living technologies for people with dementia, *Science and engineering ethics* 21 (2014) 707–765.
- [65] H. Richter, W. Mostowski, E. Poll, Fingerprinting passports, in: In NLUUG Spring Conference on Security, 2008, pp. 21–30.
- [66] J. Vongkulbhisal, Y. Zhao, R-LABS: an rfid-based indoor localisation system using antenna beam scanning, *Journal of Ambient Intelligence and Smart Environments* 5 (2013) 251–266.
- [67] A. Juels, D. Molnar, D. 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, IEEE Computer Society, Washington, DC, USA, 2005, pp. 74–88. URL: <http://dx.doi.org/10.1109/SECURECOMM.2005.59>. doi:10.1109/SECURECOMM.2005.59.
- [68] S. L. Garfinkel, A. Juels, R. Pappu, RFID Privacy: An Overview of Problems and Proposed Solutions, *IEEE Security and Privacy Magazine* 3 (2005) 34–43.
- [69] S. A. Weis, S. E. Sarma, R. L. Rivest, D. W. Engels, Security and privacy aspects of low-cost radio frequency identification systems, in: *Security in Pervasive Computing*, Springer-Verlag, 2003, pp. 201–212.
- [70] A. Juels, Rfid security and privacy: A research survey, *Journal of Selected Areas in Communication (J-SAC)* 24 (2006) 381–395.
- [71] X.-B. Li, L. 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.
- [72] J. Blocki, N. Christin, A. Datta, A. Sinha, Regret minimizing audits: A learning-theoretic basis for privacy protection, in: Proceedings of the 2011 IEEE 24th Computer Security Foundations Symposium, CSF '11, IEEE Computer Society, Washington, DC, USA, 2011, pp. 312–327. URL: <http://dx.doi.org/10.1109/CSF.2011.28>. doi:10.1109/CSF.2011.28.
- [73] M. Hildebrandt, M. Meints, RFID, profiling, and Aml, Technical Report, FIDIS (Future of Identity in the Information Society), Deliverable D7.7, <http://www.fidis.net>, 2006.
- [74] P. McDaniel, S. McLaughlin, Security and privacy challenges in the smart grid, *IEEE Security and Privacy* 7 (2009) 75–77.
- [75] W. Schreurs, M. Hildebrandt, M. Gasson, K. Warwick, Report on Actual and Possible Profiling Techniques in the Field of Ambient Intelligence, Technical Report, FIDIS (Future of Identity in the Information Society), Deliverable D7.3, <http://www.fidis.net>, 2005.
- [76] V. Christophides, V. Efthymiou, K. Stefanidis, Entity Resolution in the Web of Data, *Synthesis Lectures on the Semantic Web: Theory and Technology*, Morgan & Claypool Publishers, 2015.
- [77] P. Christen, *Data Matching - Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection*, Data-centric systems and applications, Springer, 2012.
- [78] C. Clifton, M. Kantarcioglu, A. Doan, G. Schadow, J. Vaidya, A. Elmagarmid, D. Suci, Privacy-preserving data integration and sharing, in: Proceedings of the 9th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery, ACM, 2004, pp. 19–26.
- [79] S. E. Fienberg, Privacy and confidentiality in an e-commerce world: Data mining, data warehousing, matching and disclosure limitation, *Statistical Science* 21 (2006) 143–154.
- [80] L. Sweeney, k-anonymity: A model for protecting privacy., *International Journal of Uncertainty, Fuzziness & Knowledge-Based Systems* 10 (2002) 557.
- [81] A. Narayanan, V. Shmatikov, De-anonymizing social networks, in: Proceedings of the 2009 30th IEEE Symposium on Security and Privacy, SP '09, IEEE Computer Society, Washington, DC, USA, 2009, pp. 173–187. URL: <http://dx.doi.org/10.1109/SP.2009.22>. doi:10.1109/SP.2009.22.
- [82] M. Petkovic, D. Prandi, N. Zannone, Purpose control: Did you process the data for the intended purpose?, in: W. Jonker, M. Petkovic (Eds.), *Secure Data Management*, volume 6933 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 145–168.
- [83] G. Karjoth, M. Schunter, A privacy policy model for enterprises, in: Proceedings of the 15th IEEE workshop on Computer Security Foundations, CSFW '02, IEEE Computer Society, Washington, DC, USA, 2002, pp. 271–. URL: <http://dl.acm.org/citation.cfm?id=794201.795180>.
- [84] S. Ortmann, P. Langendörfer, M. Maaser, A self-configuring privacy management architecture for pervasive systems, in: Proceedings of the Fifth ACM International Workshop on Mobility Management & Wireless Access, MOBIWAC 2007, Chania, Crete Island, Greece, October 22, 2007, 2007, pp. 184–187.
- [85] S. Resendes, P. Carreira, A. C. Santos, Conflict detection and resolution in home and building automation systems: a literature review, *Journal of Ambient Intelligence and Humanized Computing* 5 (2014) 699–715.
- [86] K. Madathil, R. Koikkara, J. Obeid, M. Geva, et al., An investigation of the efficacy of electronic consenting interfaces of research permissions management system in a hospital setting, *International journal of medical informatics* 82 (2013) 854–863.
- [87] R. R. Faden, T. L. Beauchamp, N. M. P. King, *A History and Theory of Informed Consent*, Oxford University Press, Oxford, 1986.
- [88] S. J., Informed consent, shared decision-making, and complementary and alternative medicine, *Journal of Law, Medicine and Ethics* 31 (2003) 247–250.
- [89] W. S.N., M. L.B., Physician' silent decisions: because patient autonomy does not always come first, *The American Journal of Bioethics* 7 (2007) 33–38.
- [90] B. C. M. Fung, K. Wang, R. Chen, P. S. Yu, Privacy-preserving data publishing: A survey of recent developments, *ACM Comput. Surv.* 42 (2010) 14:1–14:53.
- [91] X. Xiao, Y. Tao, Personalized privacy preservation, in: Proceedings of the 2006 ACM SIGMOD international conference on Management of data, SIGMOD '06, ACM,

- New York, NY, USA, 2006, pp. 229–240. URL: <http://doi.acm.org/10.1145/1142473.1142500>. doi:10.1145/1142473.1142500.
- [92] L. Sweeney, Achieving k-anonymity privacy protection using generalization and suppression, *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.* 10 (2002) 571–588.
- [93] P. Ohm, Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization, *Social Science Research Network Working Paper Series* (2009).
- [94] R. Bayardo, R. Agrawal, Data privacy through optimal k-anonymization, in: *Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference on, 2005*, pp. 217–228. doi:10.1109/ICDE.2005.42.
- [95] H. Park, K. Shim, Approximate algorithms for k-anonymity, in: *Proceedings of the 2007 ACM SIGMOD international conference on Management of data, SIGMOD '07, ACM, New York, NY, USA, 2007*, pp. 67–78. URL: <http://doi.acm.org/10.1145/1247480.1247490>. doi:10.1145/1247480.1247490.
- [96] A. Machanavajjhala, D. Kifer, J. Gehrke, M. Venkatasubramanian, L-diversity: Privacy beyond k-anonymity, *ACM Trans. Knowl. Discov. Data* 1 (2007).
- [97] X. Xiao, Y. Tao, Anatomy: simple and effective privacy preservation, in: *VLDB '06: Proceedings of the 32nd international conference on Very large data bases, VLDB Endowment, 2006*, pp. 139–150. URL: <http://portal.acm.org/citation.cfm?id=1182635.1164141>.
- [98] N. Li, T. Li, S. Venkatasubramanian, t-closeness: Privacy beyond k-anonymity and l-diversity, in: *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on, 2007*, pp. 106–115. doi:10.1109/ICDE.2007.367856.
- [99] C. Dwork, Differential privacy, in: M. Bugliesi, B. Preneel, V. Sassone, I. Wegener (Eds.), *ICALP (2)*, volume 4052 of *Lecture Notes in Computer Science*, Springer, 2006, pp. 1–12.
- [100] C. Dwork, A firm foundation for private data analysis, *Commun. ACM* 54 (2011) 86–95.
- [101] C. Dwork, Differential privacy: a survey of results, in: *Proceedings of the 5th international conference on Theory and applications of models of computation, TAMC'08, Springer-Verlag, Berlin, Heidelberg, 2008*, pp. 1–19. URL: <http://dl.acm.org/citation.cfm?id=1791834.1791836>.
- [102] T. Dalenius, Towards a methodology for statistical disclosure control, *Statistik Tidskrift* 15 (1977) 2–1.
- [103] M. J. Covington, W. Long, S. Srinivasan, A. K. Dev, M. Ahamad, G. D. Abowd, Securing context-aware applications using environment roles, in: *Proceedings of the sixth ACM symposium on Access control models and technologies, ACM, 2001*, pp. 10–20.
- [104] G. Chen, D. Kotz, et al., A survey of context-aware mobile computing research, Technical Report, Technical Report TR2000-381, Dept. of Computer Science, Dartmouth College, 2000.
- [105] S. Al-Rabiaah, J. Al-Muhtadi, Consec: Context-aware security framework for smart spaces, in: *Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS), 2012 Sixth International Conference on, 2012*, pp. 580–584. doi:10.1109/IMIS.2012.41.
- [106] M. Langheinrich, A Privacy Awareness System for Ubiquitous Computing Environments, in: *Proceedings of the 4th international conference on Ubiquitous Computing, UbiComp '02, Springer-Verlag, London, UK, UK, 2002*, pp. 237–245. URL: <http://dl.acm.org/citation.cfm?id=647988.741491>.
- [107] M. Nilsson, H. Lindskog, S. Fischer-Hübner, Privacy enhancements in the mobile internet, in: *IFIP WG 9.6/11.7 Working Conf. on Security and Control of IT in Society, 2001*.
- [108] G. Myles, A. Friday, N. Davies, Preserving Privacy in Environments with Location-Based Applications, *IEEE Pervasive Computing* 2 (2003) 56–64.
- [109] M. Zuidweg, J. Gonçalves, P. Filho, M. V. Sinderen, Using P3P in a web services-based context-aware application platform, in: *Proceedings of EUNICE 2003 9th Open European Summer School and IFIP WG6.3 Workshop on Next Generation Networks, 2003*, pp. 238–243.
- [110] P. Jagtap, A. Joshi, T. Finin, L. Zavala, Preserving Privacy in Context-Aware Systems, in: *Proceedings of the Fifth IEEE International Conference on Semantic Computing, IEEE Computer Society Press, 2011*.
- [111] F. L. Gandon, N. M. Sadeh, Semantic web technologies to reconcile privacy and context awareness, *Journal of Web Semantics* 1 (2004) 241–260.
- [112] A. Ranganathan, R. H. Campbell, An infrastructure for context-awareness based on first order logic, *Personal Ubiquitous Comput.* 7 (2003) 353–364.
- [113] C. S. Shankar, R. H. Campbell, A policy-based management framework for pervasive systems using axiomatized rule-actions, in: *NCA, 2005*, pp. 255–258.
- [114] B. N. Schilit, N. Adams, R. Want, Context-aware computing applications, in: *IN PROCEEDINGS OF THE WORKSHOP ON MOBILE COMPUTING SYSTEMS AND APPLICATIONS, IEEE Computer Society, 1994*, pp. 85–90.
- [115] A. Bikakis, G. Antoniou, Defeasible contextual reasoning with arguments in ambient intelligence, *IEEE Trans. Knowl. Data Eng.* 22 (2010) 1492–1506.
- [116] P. A. Bonatti, J. L. Coi, D. Olmedilla, L. 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, Springer-Verlag, Berlin, Heidelberg, 2008*, pp. 779–784. URL: http://dx.doi.org/10.1007/978-3-540-89982-2_76. doi:10.1007/978-3-540-89982-2_76.
- [117] M. F. Mokbel, C.-Y. Chow, 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, VLDB Endowment, 2006*, pp. 763–774. URL: <http://dl.acm.org/citation.cfm?id=1182635.1164193>.
- [118] N. Li, B. N. Grosz, J. Feigenbaum, Delegation logic: A logic-based approach to distributed authorization, *ACM Trans. Inf. Syst. Secur.* 6 (2003) 128–171.
- [119] P. Liu, J. bin Hu, Z. Chen, A formal language for access control policies in distributed environment, in: *Web Intelligence, 2005*, pp. 766–769.
- [120] S. Wang, Y. Zhang, Handling distributed authorization with delegation through answer set programming, *Int. J. Inf. Sec.* 6 (2007) 27–46.
- [121] G. Aucher, G. Boella, L. 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, Springer-Verlag, Berlin, Heidelberg, 2010*,

- pp. 196–213. URL: <http://dl.acm.org/citation.cfm?id=1876014.1876029>.
- [122] G. Piolle, Y. Demazeau, Representing privacy regulations with deontico-temporal operators, *Web Intelligence and Agent Systems* 9 (2011) 209–226.
- [123] M. Kanovich, P. Rowe, A. Scedrov, Collaborative planning with confidentiality, *Journal of Automated Reasoning* 46 (2011) 389–421. 10.1007/s10817-010-9190-1.
- [124] M. Blaze, J. Feigenbaum, J. Lacy, Decentralized trust management, in: *IEEE Symposium on Security and Privacy*, 1996, pp. 164–173.
- [125] M. Blaze, J. Ioannidis, A. D. Keromytis, Experience with the keynote trust management system: Applications and future directions, in: *iTrust*, 2003, pp. 284–300.
- [126] Y.-H. Chu, J. Feigenbaum, B. A. LaMacchia, P. Resnick, M. Strauss, Referee: Trust management for web applications, *Computer Networks* 29 (1997) 953–964.
- [127] D. E. Clarke, J.-E. Elie, C. M. Ellison, M. Fredette, A. Morcos, R. L. Rivest, Certificate chain discovery in spki/sdsi, *Journal of Computer Security* 9 (2001) 285–322.
- [128] K. Shilton, J. A. Burke, D. Estrin, M. Hansen, Designing the personal data stream : Enabling participatory privacy in mobile personal sensing, in: *37th Research Conference on Communication, Information and Internet Policy (TPRC)*, Arlington, VA, 2009, pp. 25–27.
- [129] A. Acquisti, The economics of personal data and the economics of privacy, *Background Paper for OECD Joint WPISP-WPIE Roundtable 1* (2010).
- [130] P. Caire, B. Alcade, L. van der Torre, C. Sombattheera, Conviviality measures, in: *10th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Taipei, Taiwan, May 2-6, 2011, 2011.
- [131] D. El Kateb, N. Zannone, A. Moawad, P. Caire, G. Nain, T. Mouelhi, Y. Le Traon, Conviviality-driven access control policy, *Requirements Engineering* (2014) 1–20.
- [132] V. Eftymiou, P. Caire, A. Bikakis, Modeling and evaluating cooperation in multi-context systems using conviviality, in: *24th Benelux Conference on Artificial Intelligence (BNAIC 2012)*, 2012, pp. 83–90.
- [133] A. Bikakis, P. Caire, Computing coalitions in multiagent systems, a contextual reasoning approach, in: *12th European Conference on Multi-Agent Systems (EUMAS 2014)*, Pragues, Dec. 18-19, 2014, 2014.
- [134] P. Hustinx, Privacy by design: delivering the promises, *Identity in the Information Society* 3 (2010) 253–255.
- [135] J. P. Hourcade, A. Cavoukian, R. J. Deibert, L. F. Cranor, I. Goldberg, Electronic privacy and surveillance, in: M. Jones, P. A. Palanque, A. Schmidt, T. Grossman (Eds.), *CHI Conference on Human Factors in Computing Systems, CHI'14, Toronto, ON, Canada - April 26 - May 01, 2014, Extended Abstracts*, ACM, 2014, pp. 1075–1080. URL: <http://doi.acm.org/10.1145/2559206.2579403>. doi:10.1145/2559206.2579403.
- [136] A. K. Elmagarmid, P. G. Ipeirotis, V. S. Verykios, Duplicate record detection: A survey, *Knowledge and Data Engineering, IEEE Transactions on* 19 (2007) 1–16.
- [137] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. M. Reagle, The platform for privacy preferences 1.0 (p3p1.0) specification, *World Wide Web Consortium, Recommendation REC-P3P-20020416*, 2002.
- [138] X. Dong, A. Halevy, J. Madhavan, Reference reconciliation in complex information spaces, in: *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, ACM, 2005, pp. 85–96.

Appendix

A. Statistics and references selection

In order to estimate the volume of scientific publications containing the term “privacy”, we did a search for the keywords “privacy” and “security” in the online repositories of the main technical publishers: IEEE Xplore, ACM Portal, Springer Online Library. Results are shown in table 3.

Publisher	Total number of publications in the database	Publications containing term “Security”	Publications containing term “Privacy”
IEEE Xplore	3,459,155	297,816	60,527
ACM Portal	2,127,336	190,400	61,322
Springer Online	N-A	268,975	62,759

Table 3: Number of publications containing the terms Security and Privacy in the major online repositories

Keyword “security” exists in around 8 % to 9 % of the entries in these databases and keyword “privacy” in around 2 % to 3 % . Roughly speaking, the volume of references containing the term privacy is around one over four of the volume of references containing the term security. To select among all these references, we did the following:

1. Select the most cited and most influential references about privacy till 2010: the references in this category are automatically selected using the help of Google Scholar and Publish or Perish tool⁸, by setting the query to return publications about privacy with more than 40 citations per year. We got 77 references in this category. Then we manually filtered out the references that are out of the scope of this study (social network privacy, cloud computing, cryptographic related), and ended up with 32 very important references as a basis for our study (these ref. are listed in table 4 in the Annex).

⁸Harzing, A.W. (2007) Publish or Perish, available from <http://www.harzing.com/pop.htm>

2. Next step is to select additional citations covering specific topics. We run new queries about privacy plus one of the following keywords: AMI, commodity, enabler, right, logic, policy, location. We manually selected the most appropriate references about these topics. We have 44 references falling in this category.
3. The recent papers (since 2011) were selected manually as we failed to find any automated way to select them. There are 36 references in this category.
4. Last category contains 26 manually selected references, not related to any specific topic and not from the most cited publications.

In order to get the statistics presented in this paper in chapters 3,4 and 5, we decided to operate on the IEEE database as a source of input for technical reasons (other publishers do not allow multiple fine-grained queries to be executed). We have exported all of the 60527 records containing the term “privacy”, and developed a classifier, which takes this set as an input and counts the number of references that fall in each category according to the keywords contained in the title, abstract and keywords fields of these records. The classifier is available online open source⁹.

⁹More info about the classifier here <https://github.com/securityandtrust/serval-classifier>

Ref.	Title	Year	Cit.	Cit./year
[96]	l-diversity: Privacy beyond k-anonymity	2007	1503	250.50
[25]	Alzheimer disease in the US population: prevalence estimates using the 2000 census	2003	1681	152.82
[114]	Context-aware computing applications	1994	3031	151.55
[80]	k-anonymity: A model for protecting privacy	2002	1728	144.00
[70]	RFID security and privacy: A research survey	2006	1143	142.88
[69]	Security and privacy aspects of low-cost radio frequency identification systems	2004	1317	131.70
[124]	Decentralized trust management	1996	2307	128.17
[98]	t-closeness: Privacy beyond k-anonymity and l-diversity	2007	818	116.86
[136]	Duplicate record detection: A survey	2007	785	112.14
[99]	Differential privacy	2006	701	87.63
[48]	Location privacy in pervasive computing	2003	920	83.64
[32]	The blocker tag: selective blocking of RFID tags for consumer privacy	2003	872	79.27
[92]	Achieving k-anonymity privacy protection using generalization and suppression	2002	944	78.67
[94]	Data privacy through optimal k-anonymization	2005	691	76.78
[90]	Privacy-preserving data publishing: A survey of recent developments	2010	297	74.25
[117]	The new Casper: query processing for location services without compromising privacy	2006	558	69.75
[43]	A taxonomy of privacy	2006	505	63.13
[55]	Unpacking privacy for a networked world	2003	628	57.09
[44]	Privacy and freedom	1970	2477	56.30
[40]	The right to privacy	1890	6955	56.09
[93]	Broken promises of privacy: Responding to the surprising failure of anonymization	2010	213	53.25
[53]	An architecture for privacy-sensitive ubiquitous computing	2004	502	50.20
[74]	Security and privacy challenges in the smart grid	2009	251	50.20
[101]	Differential privacy: A survey of results	2008	298	49.67
[97]	Anatomy: Simple and effective privacy preservation	2006	385	48.13
[18]	Privacy by design - Principles of privacy-aware ubiquitous systems	2001	621	47.77
[68]	RFID privacy: An overview of problems and proposed solutions	2005	427	47.44
[137]	The platform for privacy preferences 1.0 (P3P1.0) specification	2002	520	43.33
[138]	Reference reconciliation in complex information spaces	2005	380	42.22
[91]	Personalized privacy preservation	2006	330	41.25
[4]	Ambient intelligence: Technologies, applications, and opportunities.	2009	201	40.20
[118]	Delegation logic: A logic-based approach to distributed authorization	2003	442	40.18

Table 4: Automatically selected references