



Contents lists available at ScienceDirect

Information and Software Technology

journal homepage: www.elsevier.com/locate/infsof

Towards the development of privacy-aware systems

Paolo Guarda^a, Nicola Zannone^{b,*}^a Department of Legal Sciences, University of Trento, Italy^b Department of Information and Communication Technology, University of Trento, Italy

ARTICLE INFO

Article history:

Received 25 September 2007

Received in revised form 27 March 2008

Accepted 5 April 2008

Available online xxx

Keywords:

Privacy-related legal requirements

Requirements Engineering

Privacy policy

Privacy-Aware Access Control

ABSTRACT

Privacy and data protection are pivotal issues in nowadays society. They concern the right to prevent the dissemination of sensitive or confidential information of individuals. Many studies have been proposed on this topic from various perspectives, namely sociological, economic, legal, and technological. We have recognized the legal perspective as being the basis of all other perspectives. Actually, data protection regulations set the legal principles and requirements that must be met by organizations when processing personal data. The objective of this work is to provide a reference base for the development of methodologies tailored to design privacy-aware systems to be compliant with data protection regulations.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

In the last years, privacy has become a critical issue in the development of IT systems. This reflects the growing attention of customers to their personal data and the increasing number of statutes, directives, and regulations that are intended to safeguard the right to privacy.

The “right to privacy” was initially introduced at the end of the 19th century in the US. In 1890, Warren and Brandeis published in the Harvard Law Review an essay titled “The Right to Privacy” [1] where they defined this new right as “*the right to be let alone*”. Nowadays, there are strict regulations in place within many countries that impose rules for the collection, handling, and processing of personal data. Their main objective is to guarantee people the control on the flow of their personal data [2,3]. Organizations that handle personal data cannot escape the obligation to implement these regulations in their IT systems.

Privacy Engineering is thus emerging spurred by the realization that IT systems must comply with privacy regulations [4]. Unfortunately, it has always been difficult to bridge the gap between legal language and computer language, more importantly when legal obligations have to be converted into requirements to be enforced by IT infrastructures. Actually, invasions of privacy may not necessarily be due to malicious intents, but they were found to be inadvertently supported by technologies [5]. For instance, Kaiser Permanente (KP), a US health provider, accidentally disclosed personally identified health information (e.g., appointment details, an-

swers to patient’s questions, and medical advice) for over 800 patients through its web healthcare portal [6]. Violations of privacy may have significant consequences within an organization not only in terms of money (KP paid a \$200,000 fine [7]), but also in terms of the trust that people feel towards the organization [8].

Policies are largely used in organizations to guarantee the security of their IT systems. In the last years, many research efforts have also been devoted to safeguard the right to privacy through the use of policies. Such efforts resulted in the definition of models, languages, and standards to specify enterprises’ privacy promises (i.e., privacy policies) [9] and user preferences [10,11] as well as to enforce such promises (i.e., data protection policies) [12–14]. Such models, languages, and standards provide language constructs, but offer no methodological tools for supporting policy design.

When building IT systems that store and process personal data, designers need to define system requirements and to ensure that personal data are handled in accordance with applicable laws and regulations [15]. In most applications, informally stated and implicit privacy requirements are as urging as functional and security ones, but they are rarely analyzed and designed carefully from the beginning of the development process. Rather, they often add privacy as an afterthought, exposing the system to higher costs while endangering overall design integrity. For instance, in Europe it is in the right of data subject to decide how and for which purpose his personal data can be processed. However, it is very difficult for an organization to assure data subjects about the correct execution of data processing. It gets worse when the organization outsources data processing to an outside supplier. Even if the organization adopts proper privacy practices, because of misunderstanding of the organizational setting and the lack of standards across organizations, such practices may be pointless [16].

* Corresponding author.

E-mail addresses: paolo.guarda@unitn.it (P. Guarda), zannone@dit.unitn.it, nicola.zannone@gmail.com (N. Zannone).

It is generally accepted in the Requirements Engineering research community that system development requires models that represent the system-to-be along with its intended operational environment [17]. This is even more important when a system has to meet privacy requirements. Technologies should support proper handling of personal data within and across organizations. Moreover, privacy promises should reflect how personal data are effectively handled by the organization and available to the persons whose data are being collected. These issues emphasize the need to adequately consider privacy within a broader organizational context in order to understand the demanded system functionality and protection mechanisms. Additionally, the legal requirements constrain the technical measures, business strategies, and privacy practices of an organization.

1.1. Contribution of the paper

Our ambitious objective is to provide a reference work to assist researchers in the definition of languages and methodologies addressing the problem of developing privacy-aware systems, including the definition of privacy and data protection policies. We have thus investigated the constructs necessary to capture, represent, and analyze privacy requirements.

A first step in our endeavor is not to take such constructs for intuitively given. Only by looking at the problem from a wider legal and organizational perspective we might be able to address the problem. To this intent, we have identified the privacy principles and data subject rights that a privacy-aware system shall guarantee. Based on them, we have identified the concepts that come into play when addressing privacy concerns from a legal perspective and how these concepts have been interpreted from a technological perspective. We have founded our work on the European Directive on data protection (EU Directive 95/46/EC). To broaden the audience, we have also compared the EU legal system with the US one, pointing out similarities and differences of such systems. This work thus serves also as a bridge between computer scientists and legal experts in order to facilitate interactions between them.

The analysis of privacy principles set by the data protection regulations has also revealed the need of understanding the organizational setting in which a privacy-aware IT system operates. For instance, different measures should be taken by the data controller in case he assigns the processing of personal data to an employee within the same organization or in the case the data processing is outsourced to an external recipient. Requirements Engineering can aid system designers by providing the tools necessary to model and analyze the organizational context of a system in terms of the structure and goals of an organization. We have thus studied the current proposals in this research area and, in particular, those that explicitly address privacy concerns. In this study we have also analyzed proposals intended to assist organizations in verifying and guaranteeing the consistency among enterprise goals, privacy policies, and data protection policies.

The paper is structured as follows. Next section presents an overview of data protection regulations in Europe along with the privacy principles and data subject rights established by such regulations. Section 3 presents privacy-related concepts and places them in the EU legal framework. Section 4 introduces the problem of Privacy Engineering. Sections 5 and 6 review the state of the art on languages for specifying privacy policies and user preferences and languages for specifying data protection policies, respectively. Similarly, Section 7 presents a critical review of the existing proposals for Privacy Requirements Engineering. Section 8 discusses the alignment among privacy artifacts and provides an overview of existing proposals that attempt to guarantee their consistency.

Section 9 compares the EU legal framework with the US legal framework, showing similarities and differences between them. Finally, Section 10 concludes the paper with final remarks.

2. Privacy and Data Protection in Europe

In the '50, the Council of Europe recognized privacy as a fundamental right. This right was defined in article 8 of the European Convention of Human Rights and Fundamental Freedoms (Council of Europe, Rome, 1950), which establishes that everyone has the right to respect for his private and family life, his home and his correspondence [18].

Though the right to privacy was vested in Europe since 1950, privacy regulations were introduced in European countries legislation only much later. The first cases of legal intervention were in West Germany with the statutes of Assia (October 7th, 1970) and Bavaria (October 12th, 1970), and then a federal statute on data protection (Bundesdatenschutzgesetz, BdsG) in 1977. National statutes were also issued in Sweden (1973), France (1978), Luxembourg (1979), Denmark (1979), Austria (1980), Norway (1980), Iceland (1982), United Kingdom (1984), Finland (1988), the Netherlands (1990), Portugal (1991), Spain (1993), Belgium (1993), and Switzerland (1993). Moreover, Spain, Portugal, Austria, the Netherlands, Germany and Greece have amended their own Constitution including privacy clauses [19].

The European Union recognized the need to harmonize data protection laws across Europe in order to achieve two main objectives, namely the protection of citizens' privacy and the maintenance of free flows of personal data across Member States [2,3]. To this end, the European Union has enacted several acts with the objective of regulating the management of personal data. These acts also intended to regulate the computer-supported processing of personal data. Table 1 gives a brief history of EU Directives on data protection [2,20].

It is worth noting that EU Directives only set general principles and leave each Member State to implement specific national measures. For instance, the Italian statute on data protection (December 31st, 1996, n. 675) was issued to implement the Directive 95/46/EC (hereafter the "Directive"). It was successively replaced by the Italian Data Protection Code (June 30th, 2003, n. 196) (hereafter the "Code"), which gathers up all the old Italian acts on this field and gives new rules in a systematic and organic way.

2.1. Privacy Principles

Data protection regulations in the EU set the main principles that establish how data processing shall be performed. Legal experts refined these general principles, adding new principles or making explicit some others [2,3]. We can summarize privacy principles as follows:

Table 1
Brief history of the European data protection legislation

Directive 95/46/EC of the European Parliament and the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data
Directive 97/66/EC of the European Parliament and of the Council of 15 December 1997 concerning the processing of personal data and the protection of privacy in the telecommunications sector
Directive 2002/58/EC of the European Parliament and the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications), that abrogated the directive of 1997

- (1) *Fair and Lawful Processing*: the collection and processing of personal data shall neither unreasonably intrude upon the data subjects' privacy nor unreasonably interfere with their autonomy and integrity, and shall be compliant with the overall legal framework.
- (2) *Consent*: personal data shall be collected and processed only if the data subject has given his explicit consent to their processing.
- (3) *Purpose Specification*: personal data shall be collected for specified, lawful and legitimate purposes and not processed in ways that are incompatible with the purposes for which data have been collected.
- (4) *Minimality*:¹ the collection and processing of personal data shall be limited to the minimum necessary for achieving the specific purpose. This includes that personal data shall be retained only for the time necessary to achieve the specific purpose.
- (5) *Minimal Disclosure*: the disclosure of personal data to third parties shall be restricted and only occur upon certain conditions.
- (6) *Information Quality*: personal data shall be accurate, relevant, and complete with respect to the purposes for which they are collected and processed.
- (7) *Data Subject Control*: the data subject shall be able to check and influence the processing of his personal data.
- (8) *Sensitivity*: the processing of personal data, which are particularly sensitive for the data subject, shall be subject to more stringent protection measures than other personal data.²
- (9) *Information Security*: personal data shall be processed in a way that guarantees a level of security appropriate to the risks presented by the processing and the nature of the data.³

2.2. Data Subject Rights

Data protection regulations also set data subject rights. For instance, according to the EU Directive (§12 et seq.), data subject has:

- the right of access to his personal data;
- the right to object to a data processing;
- the right to delete his personal data;
- the right to have inaccurate personal data updated or deleted;
- the right to prevent that personal data are used to achieve purposes different from which he has given the consent.

3. The basic concepts

In this section, we define the concepts necessary to understand and express privacy-related legal requirements. Such definitions are based on the EU Directive in the light of the Italian Code. They are general enough to cover the entire spectrum of socio-technical systems, from organization procedures to IT solutions. We also

¹ This principle covers the “data minimization” principle defined by the Italian Code (§3) as well as the “least privilege” principle proposed by Saltzer and Schroeder [21]. Other terms, such as “need-to-know”, “necessity”, “non-excessiveness”, “proportionality”, “frugality” are used to refer to this principle [2].

² Many legal systems requires organizations to notify the privacy authority before processing sensitive data.

³ The EU Directive (§17) requires data controllers to implement security measures for ensuring that personal data are protected from accidental and unlawful destruction, alteration or disclosure; such measures have to be commensurate with the risks involved in the data processing having regard to the state of the art and the cost of their implementation. In the Italian Code (§31), there is no reference to the implementation cost. This makes Italian regulations much more restrictive.

compare the terminology and contrasting definitions used by computer scientists and legal experts in order to bridge the gap between them and so facilitate their interaction.

3.1. Categories and typologies of data

Data play a fundamental role in privacy domain as they shall be collected, processed, and disclosed according to the privacy principles defined in Section 2.1. Different kinds of data can be involved in a processing:

- *Personal data*: any data that can be used to identify a person⁴ (EU Directive §2, lett. a, and Italian Code §4, co. 1, lett. b);
- *Sensitive data*: any data that disclose information about racial or ethnic origin, religious, philosophical or other beliefs, political opinions, membership of parties, trade unions, associations or organizations of a religious, philosophical, political or trade-unionist character, as well as personal data disclosing health and sex life. An important subcategory of this kind of data are medical data (EU Directive §8, and Italian Code §4, co. 1, lett. d);
- *Identification data*: personal data that permit the direct identification of the data subject (Italian Code §4, co. 1, lett. c);
- *Anonymous data*: any data that cannot be associated to any identified or identifiable data subject (Italian Code §4, co. 1, lett. n). This category of data is not regulated by data protection regulations.

The distinction of categories of data is necessary for the principles of sensitivity and information security since the measures adopted to protect data shall be adequate to the nature of data. As a concrete example, we mention the “Documento Programmatico sulla Sicurezza”⁵ issued by the University of Trento, which establishes that university staff shall change their password every six months, but employees accessing sensitive data shall change their password every three months.

3.2. Actors

Different actors can be involved in a data processing. The EU Directive and the Italian Code identify the following actors:

- *Data subject*:⁶ the person to whom personal data refer (Italian Code §4, co. 1, lett. 1).
- *Data controller*: the person who determines the purposes for which and the manner in which personal data are processed (EU Directive §2, lett. d, and Italian Code §4, co. 1, lett. f).
- *Data processor*: any person who processes personal data on behalf of the data controller (EU Directive §2, lett. e, and Italian Code §4, co. 1, lett. g).⁷

⁴ Differently from the EU Directive, the Italian Code also refers to “legal” person, besides “natural person”.

⁵ The “Documento Programmatico sulla Sicurezza” (DPS) is a document founded on the ISO/IEC 17799 standard [22], which describes the security and privacy policies adopted by an organization. The Italian Code enforces public and private organization to issue the DPS annually. This document should describe the analysis of risks affecting personal data as well as the measures adopted by the organization to protect them from possible abuse.

⁶ The concept of data subject is also expressed using the terms *donor of the personal information* [23] or *data owner* [15]. However, they are not equivalent in the EU legal framework. For instance, the latter relates privacy to the concept of property. On the contrary, privacy is a fundamental right in the EU legal framework.

⁷ In the EU legal framework, the relationship between the controller and the processor must be governed by a contract or a legal agreement. Due to the nature of this relationship, a data processor cannot be an employee of the data controller [3]. On the contrary, in the Italian legal context it could be also a member of the organization of the data processor.

- *Persons in charge of the processing*: any person that has been authorized by the data controller or processor to carry out processing operations (Italian Code §4, co. 1, lett. h).⁸
- *Third party*: any person than the data subject, controller, processors, and persons in charge of the processing (EU Directive §2, lett. f).
- *Recipient*: any person to whom data are disclosed, whether a third party or not (EU Directive §2, lett. g).
- *Privacy authority*: special authorities appointed to oversee the implementation of the data protection laws (EU Directive §28, and Italian Code §153 et seq.).⁹

The identification of the actors involved in the data processing is necessary to set the responsibilities and powers imposed by the privacy principles. If we consider, for instance, the organizational structure of universities, the above actors can be identified in the following subjects: the data controller is the chancellor as the legal representative of the university; the data processors are the administrative managers, faculty deans, and heads of department as well as external suppliers to whom the university has outsourced data processing; the persons in charge of the processing are the university staff appointed by the data controller and processors to carry on data processing; data subjects are students, professors, employees, etc.

3.3. Purpose

The *purpose* is the rationale of the processing, on the basis of which all the actions and treatments have to be performed.

The purpose specifies the reason for which data can be collected and processed. Essentially, the purpose establishes the actual boundaries of data processing. The notion of purpose plays a key role in data protection and it is at the basis of most of the principles presented in Section 2.1. Here we recall the purpose specification principle according to which an organization can collect personal data for specified, lawful and legitimate purposes. Any other kind of processing is not allowed, unless explicitly permitted by the data subject.

To simplify the management, purposes can be organized in a hierarchical structure [24]. This structure should support the specification of privacy policies and data protection policies that govern organizations' business strategies. Organizations generally provide their services in different ways. Organizations might also need to decompose a generic purpose into more specific ones since they are not completely able to provide demanded services by themselves. This is the case for a business process where different partners explicitly combine their efforts into one process in order to provide a service to customers [25].

As a partial solution, Agrawal et al. [23] proposed to decompose purposes into multiple sub-purposes and then store them in the database. However, using this simple notion of sub-purpose we lose the logical relation between a purpose and its sub-purposes. Consequently, it does not allow for the reasoning about the fulfillment of root purposes. For example, a customer might opt out of providing information necessary to fulfill a sub-purpose that, however, is necessary to fulfill the root purpose. Thereby, the organization collects from the customer information that is altogether insufficient to fulfill the root purpose, breaking minimality and information quality principles.

Another solution is proposed by Karjoth et al. [26] who consider purposes as strings that identify the intentions for which an oper-

ation can be executed. In their approach, purposes are ordered in a hierarchical manner with a directory-like notation. In this setting, if an operation is allowed for a given purpose, it is also allowed for all sub-purposes. Similarly, Byun et al. [14] organized purposes according to a hierarchical structure based on the principles of generalization and specialization. Yet, these approaches do not support the specification of alternatives, thereby limiting reasoning about the fulfillment of root purposes.

There is evidence that the goal-oriented approach is adequate to model complex business strategies [27,28]. Based on such an approach, Massacci et al. [29] have organized purposes into AND/OR trees. A similar approach was proposed in [25] where hypergraphs are used to represent purpose hierarchies. Here, AND/OR-decompositions are represented as hyperedges.

3.4. Consent

The *consent* is a unilateral action producing effects upon receipt that manifests the data subject's volition to allow the data controller to process his data.

According to the EU Directive (§2, lett. h) and the Italian Code (§23), processing of personal data by private entities or profit-seeking public bodies shall be allowed only if the data subject gives his/her explicit consent. This corresponds to the Principle of Consent.

Different solutions have been proposed to model the notion of consent. Several approaches (e.g., [30–32]) identify consent with the notion of permission and use it to model the ability to perform actions in a system. Other proposals represent consent as a precondition for delegation [33].

Data subjects can withdraw the consent at any time exercising the rights that the data protection laws recognize to them, as the right to object to the processing or to delete collected data (EU Directive §14 and Italian Code §7). As consequences, a privacy-aware infrastructure shall allow data subjects to withdraw their consent. Different models for permission revocation have been proposed [34–38], but no all of them may be adequate to model the withdrawal of data subject's consent. For instance, in [36] two models for permission removal are presented, namely *deep removal* and *shallow removal*. The deep removal of a permission results in recursively removing all privileges that are consequence of it. On the contrary, backward propagation of permission removal is not used in shallow removal. Therefore, a deep removal approach is the more appropriate for our purposes since the withdrawal of the consent by the data subject to the data controller implies that the data processors and persons in charge of the processing appointed by that controller lose the authorization to process data as well.

As final remark we want to point out that though the consent may be intuitively seen as a contract, the right of data subjects to withdraw it and the inalienability of fundamental rights, as privacy is, make a contractual approach inadequate to data protection in the European legal system. This approach however can be adopted in other legal systems [39,40].

3.5. Obligation

An *obligation* is a condition or an action that is to be performed before or after a decision is made.

Although obligations are not explicitly mentioned in the privacy principles of Section 2.1, they provide a means for their implementation. This is, for instance, the case of the Principle of Minimality where obligations can be used to delete data once the retention period associated with them is expired. Obligations can also help in implementing the Principle of Information Security as they impose constraints on how the data may be used [15].

⁸ This actor is not clearly defined in the EU Directive.

⁹ Most countries with data protection laws have established these special authorities. In carrying out their tasks, they are required to be functionally independent of the governments and/or legislatures which establish them. The powers of data protection authorities are often broad and largely discretionary. In most cases, they are empowered to issue legally binding orders.

The term obligation was initially used by Damianou et al. [41], who introduced the notion of obligation policy for specifying the actions that must be performed when certain events occur. Successively, it has been recognized that access decisions cannot depend only on the identity and authorization of the entity requesting the access, but it is also necessary to consider the consequences of an access [42]. The concept of obligation was thus integrated in access control frameworks (e.g., [12,42–44]) to define the actions that must be taken when permissions are granted. Obligations are also used to restrict the set of permissible actions (i.e., actions that do not comply with obligations are not permitted) [13] and to specify the intended usage of data [15]. A characteristic of these approaches is that obligations are incomparable with authorizations. Essentially, obligations are conceptually distinguished by authorizations. A different approach is adopted in Deontic Logic [45] where obligations entail permissions. Finally, we mention speech act [46] where obligations are related to promises.

Obligations can also be classified according to their nature. For instance, Hilty et al. [15] proposed a classification of along two dimensions, namely time and distribution. In the temporal dimension obligations are distinguished with respect to bounded time and invariance properties. In the distribution dimension obligations are classified according to their observability nature. Understanding the nature of obligations will help designers in the selection of suitable enforcement mechanisms.

The term obligation is not consistently used in the literature. For instance, in [47] it refers to the actions that have to be performed before or during a usage exercise. Some authors [43,15] have distinguished between actions that must be performed before an access is authorized, called *provisions*, and actions that must be performed after an access is authorized, called *obligations*. The term obligation is used with a different meaning in the legal language. Here it indicates “a legal agreement specifying a payment or action and the penalty for failure to comply”. In particular, in Civil law the term “obligation” refers to legal agreement in general, or to the fulfillment. What we call obligations refer to “accessory (or ancillary) obligations” that are secondary duties involved in an agreement. In Common law, the term “obligation” refers to a pathological situation of responsibility (e.g., breach of a contract) that imposes a sanction. Therefore, it is related to the notion of “liability”.

Coupled with the notion of obligation, we can find the notion of *compensation action* [15,43]. Essentially, compensations actions are actions that are taken when obligations are not fulfilled.

3.6. Retention Period

The *retention period* defines how long data shall be kept.

Retention period is inevitably related to the Principle of Minimality that requires the data controller to delete, destroy, or anonymize¹⁰ personal data when the processing purpose is fulfilled [3]. Many access control frameworks [48–50] provide support for revoking authorizations when the associated temporal interval expires. However, the issue of deleting data is still challenging. For instance, personal data shall not only be deleted from the database, but also from the logs without affecting recovery [23].

It is worth noting that the notion of retention period is different from *data retention*. Data retention refers to the storage of call detail records and Internet traffic and transaction data by governments and commercial organizations. It is related to public security issues and to oppose the criminality (Directive 2006/24/EC on the retention of data generated or processed in connection

with the provision of publicly available electronic communications services or of public communications networks).

4. Privacy Engineering

Privacy Engineering was defined by Kenny and Borking “as a systematic effort to embed privacy relevant legal primitives into technical and governance design” [4]. Privacy Engineering is thus the discipline addressing the development of models, tools, processes, and methods needed to design systems that guarantee privacy protection according to privacy regulations.

Privacy Engineering has been tackled by several research communities, such as Privacy Requirements Engineering [51–53], Privacy Policy and User Preference Specification [9–11], Privacy-Aware Access Control [12–14], Identity Management [54–56], Digital Rights Management [57], etc. Even though these fields are all important in the development of privacy-aware systems, we choose to only focus on Privacy Requirements Engineering, Privacy Policy and User Preference Specification, and Privacy-Aware Access Control as we see these fields as key topics for this article's intentions. Indeed, we are interested in methods that assist organizations in the specification of privacy promises, guaranteeing their enforcement, and – last but not least – ensuring their compliance with data protection legislation.

Several contributions addressing privacy can be found in the area of policy specification. They provide language constructs for representing privacy requirements, but do not offer methodological tools for supporting organizations in the design of their policies. In this setting, the inclusion of privacy is usually done on top of the existing design. This is a critical issue since privacy features have to be fitted into a system that might be not able to accommodate them. What is needed is a methodology for describing organizations and their operational procedures, and then deriving privacy policies and protection mechanisms from them. For instance, privacy promises are used by organizations to get the informed consent from data subjects. Without an understanding of the organizational structure and procedures, such promises might not correspond to the actual practices of the organization. Yet, the organization shall guarantee that privacy promises are enforced properly. For instance, the principle of privacy specification requires that data cannot be processed in a way that is incompatible with the purpose for which the data have been collected. These issues demand organizations to guarantee the consistency among privacy artifacts (i.e., enterprise goals, privacy policies, user preferences, and data protection policies).

Requirements Engineering seems to have the potentialities to help policy designers in their task. Actually, Requirements Engineering offers methods to analyze not only the system-to-be but also its operational environment, leading to the full comprehension of organization practices. Moreover, Requirements Engineering enables the representation of the system and its operational environment at different abstraction levels, providing a means for reconciling privacy promises and their enforcements.

In the remainder of the paper, we have investigated existing specification languages for privacy policy and Privacy-Aware Access Control (Sections 5 and 6, respectively) as well as the features that Requirements Engineering methodologies should provide to system designers (Section 7). Another important aspect, which we are interested in, is to guarantee the consistency among privacy artifacts. For this purpose, we have analyzed the relationships among the privacy artifacts produced during the different phases of the system development process (Section 8).

A difficulty we met in this work was due to the terminology that is often used in the literature. Many authors have used the term “privacy policy” to indicate both statements describing privacy promises and statements specifying their enforcement [58,52,26].

¹⁰ The anonymization process consists in removing personal identifiers. Once all personal identifiers are removed, the information ceases to identify an individual and so it ceases to be personal data.

However, these types of statement address different problems and, consequently, they should specify different information and at a different level of granularity. To overcome this difficulty, we have clearly distinguished them in this paper. Specifically, we have borrowed the term *data protection policy* from [15] to indicate statements that specify how enterprise promises are enforcement within the enterprise. On the other hand, the term *privacy policy* is used in the paper only to refer to the policies disclosed by organizations to inform data subjects about how personal data are being managed and their rights.

5. Privacy Policies and User Preferences

Organizations must inform data subjects about the use of their personal data (EU Directive §10 and §11, and Italian Code §13). This measure is necessary to comply with the privacy principles stated in Section 2.1. For instance, organizations need to get informed consent from data subjects before processing personal data (the Principle of Consent). Yet, data subjects shall be able to check and influence the processing of their data (the Principle of Data Subject Control). Without a knowledge of organization's privacy practices, data subjects cannot wield their rights. The disclosure of privacy practices is also relevant for the Principle of Purpose Specification as organizations cannot process personal data in a way incompatible with the purpose for which such data have been collected.

Usually, organizations achieve the obligation of informing data subjects by means of privacy policies. A *privacy policy* is essentially a comprehensive and high-level description of organization's privacy practices [59]. Privacy policies shall contain information regarding:

- the purposes and modalities of the processing for which the data are collected;
- the obligatory or voluntary nature of providing the requested data;
- the categories and typologies of data concerned;
- the recipients or categories of recipients;
- the rights of data subjects.

Table 2 shows the skeleton of the privacy policy defined by the Italian privacy authority for its portal. We refer to [60] for the entire informative note.

It is worth noting how the Italian privacy authority has considered and specified the use of Privacy Enhancing Technologies and, in particular, P3P in the privacy policy of its portal (see item 7 in Table 2). P3P (Platform for Privacy Preferences) [9] is an emerging W3C standard for the specification of privacy policies. Its goal is to enable users to gain more control over the use of their personal data on web sites they visit. P3P enables web sites to express their data-collection and data-use practices in a machine-readable XML format that can be retrieved automatically and interpreted easily by users. P3P policies enumerate the collected data, explain how those data will be processed, and specify the purpose for which

they will be collected and processed, and for how long it will be stored (i.e., retention period). The purpose may also have an attribute that allows data subjects to express their consent to the purpose. In addition, P3P policies specify the data recipients and other information, for instance, about dispute resolution and the address of a site's human-readable privacy policy. In summary, P3P provide websites with the support for describing their privacy promises.

The designers of P3P have simultaneously designed P3P Preference Language (APPEL) [11]. An APPEL preference is a set of rules that consist of a judgment, expressed in terms of block, limited, and request, and a condition under which the judgment is issued. Ideally, the use of P3P and APPEL allows a data subject to check a privacy policy against his privacy preferences and then automatically determine if his personal data can be disclosed. To support the negotiation between organizations and customers, Agrawal et al. [61] proposed a server-centric architecture for P3P. The P3P protocol has two parts: *Privacy Policies*, in which a web site can encode its data-collection and data-use practices using P3P, and *Privacy Preferences*, in which customers can specify their privacy preferences using APPEL. Database querying is then used for matching user privacy preferences against privacy policies.

Though P3P and APPEL are used in real applications (e.g., Internet Explorer 7), many drawbacks of these languages have been noticed [10,62–64]. For instance, APPEL allows users to specify what is unacceptable in a policy, but not what is acceptable [10]. Hogben [62] recognized the limitations of P3P in cookie management, user interfaces and vocabularies, and the ambiguity and awkwardness of APPEL. Schunter et al. [63] showed the ambiguity of P3P and argued that this language does not provide clear guidelines for policy design and interpretation. Yu et al. [64] identified the lack of a well-defined semantics for P3P policies. Another criticism concerning P3P is the current inability of P3P vocabulary to accurately replicate human-readable policies [18]. This drawback creates legal uncertainty that affects the legal value of P3P statements, especially for what concerns liability issues [65].

To overcome the limitations of APPEL, Agrawal et al. [10] proposed XPref, an XPath-based privacy preference language. XPref has many advantages over APPEL in terms of clarity, ease of use, and expressiveness. However, XPref, as well as P3P, is still a syntax-based preference language and, thus, it does not solve all APPEL's problems. Based on this observation, Yu et al. [64] defined a formal semantics for P3P policies, which precisely identifies the relationships between the components of P3P statements (i.e., data, purposes, recipients, and retention period). On the basis of such a semantics, they have proposed SemPref [66], a semantics-based preference language for P3P.

Another framework for the specification of privacy policies and user preferences has been proposed by Tumer et al. [67]. Here, enterprises specify which information is mandatory for achieving a service and which is optional, while customers specify the type of access for each part of their personal data: free (i.e., the access is granted without conditions), limited (i.e., the access is granted only if the enterprise has defined as mandatory that part of information), or not given (i.e., the access is never granted). Then, the framework matches enterprise policies with customer preferences. If mandatory information is not given by a customer, the framework verifies if alternative strategies stated by the enterprise match customer preferences in order to reach an agreement with the customer.

Languages for privacy policy specification and, in particular, P3P are expressive enough to represent organization privacy practices as demanded by privacy regulations. In other words, the XML elements that form a P3P policy allows policy writers to specify the information set by privacy regulations. Nonetheless, P3P does not prevent the specification of unfair and deceptive practices. It gets even worst when user agents enter in the picture. For instance,

Table 2
Skeleton of the privacy policy of Italian privacy authority portal

1. Personal Data Controller
2. Place of personal data processing
3. Categories and typologies of personal data to be processed
3.1. Navigation data
3.2. Data supplied by users on a voluntary basis
3.3. Use of Cookies or users tracing/monitoring systems
4. Optional nature of data supplying
5. Data processing modalities
6. Data subject rights
7. P3P

P3P user agents are software agents designed to fetch P3P policies, interpret them, display them, and perform actions based on policies and user preferences. Such agents usually present P3P policies in a simplified manner, for instance, by replacing technical terms with more informal descriptions. This simplification process intends to increase usability and readability of policies, but reduces the precision of policy terms. Such a lack of precision may raise legal implications for users, organizations, and user agent developers. For instance, misunderstanding of privacy policy terms may make privacy agreements invalid. Thereby, together with the definition of policy languages it is necessary to develop infrastructures and applications that address privacy issues properly. A discussion on privacy-aware infrastructures and applications, however, is out of the scope of the paper.

6. Privacy-Aware Access Control

In the previous section, we have reviewed existing languages for the specification of privacy promises. Organizations shall adopt the measures necessary to enforce such promises. Actually, the Principle of Purpose Specification requires that data processing is consistent with the purposes for which personal data have been collected. In addition, the Principle of Information Security enforces organizations to guarantee a level of security appropriate to the risks presented by the processing and the nature of the data. Data protection policies have been proposed to address those issues.

A typical access control policy is expressed as a tuple $\langle s, o, a \rangle$ with the intended meaning that a subject s can perform an action a on an object o [68–71]. These three elements, however, are insufficient to specify data protection policies [52,29]. In addition to the above three basic authorization elements (subjects, objects, and actions), Karjoth et al. [72] identified other three elements that shall occur in data protection policies, namely purpose, condition, and obligation. Based on this observation, a number of languages and models tailored to specify and enforce data protection policies were proposed [58,13,14,26,44]. Their aim is to support enterprises in keeping the promises made to customers.

The most prominent proposals are E-P3P [26] and its successor EPAL [58,73]. E-P3P (Platform for Enterprise Privacy Practices) has been proposed to enable an enterprise to formalize the exact privacy policy that shall be enforced within the enterprise. It formalizes the privacy promises into data protection policies and associates them to each piece of collected data they refer. These sticky policies are then used in access control decisions to enforce the privacy promises made. The E-P3P policy language categorizes the data an enterprise can collect and the rules which govern the usage of these data. An E-P3P policy is essentially a set of privacy rules that define users, actions, data, purpose, conditions, and obligations. Part of these elements (e.g., user, data, and purpose) are also structured hierarchically along the lines given by Jajodia et al. [70] to ease policy design and management as well as to understand changes in policy specifications. Based on E-P3P, EPAL (Enterprise Privacy Authorization Language) was proposed by IBM as part of its enterprise privacy management solution. EPAL defines an XML-based syntax to formulate privacy practices for enterprise-internal enforcement.

Close to EPAL, we can find eXtensible Access Control Markup Language (XACML) [44]. XACML is an OASIS standard for access control. This standard provides a policy model used to describe access control policies. In addition, XACML has a request/response language to express queries about whether or not a given permission should be allowed to a certain user. XACML and EPAL are very similar in concept, though XACML does not have a special construct for specifying purpose, like EPAL. However, it has been added in XACML's privacy policy profile [74]. A comprehensive comparison between XACML and EPAL was given by Anderson [75].

Many other proposals can be found in the literature. For instance, Barth et al. [13] proposed Declarative Privacy Authorization Language (DPAL). Differently from previous privacy languages, DPAL supports the perspectives of both enterprises and customers. When interpreting a DPAL policy, each rule in the policy is enforced, enabling combination, unlike in EPAL. However, DPAL does not ensure policy consistency. It is assumed the existence of an algorithm able to detect inconsistencies in DPAL policies, but the details of such an algorithm were not given nor its effectiveness and efficiency were discussed. Byun et al. [14] proposed a purpose-based access control framework extending RBAC [71] along the lines given by Agrawal et al. [23] in their Hippocratic database systems. The aim of this framework is to enforce privacy promises encoded in privacy policy languages, such as P3P, in database management systems. The framework is based on the notion of intended purpose, which specifies the intended usage of data, and the notion of access purpose, which specifies the purpose for which a given data element is accessed. They also introduce purpose hierarchies and a purpose management model for reasoning on access control. Another important issue addressed in this work is the data labeling scheme that specifies how data are associated with intended purposes. Based on this labeling scheme, a database system only returns the data that can be accessed for given purposes. This approach, however, solves only one part of the data protection problem: it controls who can access which data for which purpose, but not how the data are used once accessed.

Above frameworks provide a formalism to specify data protection policies, which, in most cases, is able to deal with the requirements set by data protection regulations. Together with a policy language, they also provide methods for evaluating and enforcing policies. The problem of those methods is that they have been built to manage policies within single organizations rather than in a distributed system. Nowadays outsourcing is a common business practice adopted by public and private organizations to reduce costs. Outsourcing has, however, a strong impact on the data protection requirements of organizations as personal data are disclosed to an external supplier over whom the data controller may not have direct control. These issues are partially addressed by Mazzoleni et al. [76] who proposed to extend XACML with algorithms addressing the problem of determining policy similarities and of policy integration across autonomous organizations. However, they do not consider obligations, which are a fundamental ingredient to enforce data protection requirements. Hilty et al. [15] proposed to use Distributed Temporal Logic to formalize data protection policies. Such policies are defined in terms of authorizations and obligations. Based on the nature of obligations, they provide strategies for enforcing obligations in distributed systems. The drawback of this proposal is that obligations are also used to represent the intended use of data, making it difficult the policy management.

7. Privacy Requirements Engineering

Requirement Engineering is “the branch of software engineering concerned with the real-world goals for, functions of, and constraints on software systems. It also concerned with the relationship of these factors to precise specifications of software behavior and to their evolution over time and across software families.” [77]. Requirement Engineering is particularly critical, because misunderstandings in this phase of the development process may lead to expensive errors in the deployed system. In this section, we propose the features necessary for a Requirements Engineering framework addressing the problem of developing privacy-aware systems.

A Privacy Requirements Engineering methodology should provide a specification language for representing privacy and data protection requirements in the organization domain as well as systematic methods for eliciting and analyzing these requirements. A requirements specification language should consist of a set of primitive constructs that allow one to express and relate the notions proposed in Section 3. From a methodological perspective, the framework should comprise the activities to capture, represent and analyze privacy requirements along functional and security requirements. Accordingly, besides the traditional activities provided by Requirements Engineering frameworks, we have identified the following activities:

- capture the structure of organizations and their environmental setting by identifying the different actors defined by the privacy regulations;
- capture the purposes for which personal data are collected and link permissions to them;
- identify the kind of data involved in the processing;
- capture the obligations that shall be fulfilled by an actor and link them to the permission that has generate them.

In the last years it has been recognized the importance of capturing and modeling privacy requirements in the early stages of system development to provide high assurance of privacy protection to both organizations and their customers [52]. This has spurred several researchers to use and extend existing Requirements Engineering frameworks to cope with security and privacy issues. Most frameworks address privacy along with other security requirements, rather than as a separate design criterion in the system development process [78]. As consequence, they fail to capture the whole range of privacy-related legal requirements. For instance, in [77] security and privacy are considered as vague goals to be satisfied, while a precise description and enumeration of specific security and privacy properties and behavior are missing. Similarly, the Non-Functional Requirements (NFR) framework [79] treats security and privacy requirements as non-functional (or quality) requirements and models them as softgoals, that is, goals having no clear-cut definition of their satisfaction. Liu et al. [80] have extended this approach by offering facilities for threats, vulnerabilities and countermeasures analysis. Though softgoals allow for an explicit identification and evaluation of alternative ways by which stakeholders can achieve their goals [81], they fail to capture privacy aspects. For instance, these approaches do not support the notion of purpose, which is a central for understanding most privacy concerns.

Moving towards this direction, Kaindl [82] proposed a systematic design process based on a model combining scenarios with goals and functions. In this combined model, purpose serves as a link between functions and goals: system functions have some purposes and these purposes match the goals of the users. However, this framework lacks the concept of permission and the link with the purpose for which permission has been granted, which are necessary to deal with data protection.

This issue is addressed in Goal-Based Requirements Analysis Model (GBRAM) [83] and, especially, in its extensions [51,84,52]. GBRAM provides a methodological approach to identify system and enterprise goals as well as requirements. In particular, the framework provides heuristics for identifying the goals that systems must achieve, managing trade-offs among them, and refining them into operational requirements. This work has been successively extended to derive privacy policies [51,84] as well as data protection policies [52] from organizational goals. These extensions have been discussed in details in Section 8. A drawback of GBRAM as well as its extensions is that they do not provide facilities for a systematic analysis of the organization structure. Therefore, it is not able to capture the different nature of obligations

that actors must fulfill due to the role played in the privacy context.

The importance of linking permission to purpose was also recognized by Massacci et al. [53], who introduced the notion of *special power of attorney* for which an appointed attorney is vested solely with the power needed to carry out a specific affair. Along this direction, they extended Secure Tropos [31], an agent-oriented security requirements engineering methodology, to address the analysis of privacy requirements. Their original proposal employs the notions of supervision, permission, delegation, and trust to capture organizational and security aspects of socio-technical systems and provides requirements engineers with formal analysis techniques for requirements verification as well as for the verification of the consistency between security, privacy, and functional requirements. In particular, it allows for the verification of the need-to-know principle by ensuring that actors have permission only if they actually do need such permission. In [29], the authors have refined the notions of permission and delegation offered by Secure Tropos by making explicit the purpose for which permission is granted. Though this framework also allows for the modeling of the obligations an actor shall fulfill in terms of the tasks he has to execute, it does not provide support to link them to the permission that has demanded them.

The analysis of existing Requirements Engineering methodologies has shown that they are not sufficient to address privacy and data protection issues. For instance, some proposals do not allow the specification of fundamental notions such as purpose and permission, whereas other proposals lack activities for the elicitation, modeling and analysis of certain aspects necessary to understand the privacy domain such as the analysis of organization structures.

8. Aligning privacy artifacts

Ensuring the correctness of privacy-aware systems also requires the alignment and compliance among privacy artifacts (i.e., enterprise goals, privacy policy, user preferences, and data protection policies) introduced in the system development process. In the remainder of this section, we discuss the relationships between privacy artifacts and review proposals addressing this issue.

8.1. Aligning enterprise goals and privacy policies

Defining privacy policies and bringing them into alignment with the organizational setting are complex activities. These activities require one to understand what are the organizational goals, the structure of the organization and its environmental setting. Their failure makes organizational goals, privacy policies, and system requirements to be misaligned and, consequently, it is extremely difficult for system designers to demonstrate that their systems are privacy-aware. Conversely, organization practices should be reflected in the actual privacy policies and system requirements, and privacy policies in the actual system requirements [51].

This awareness has been matched by a number of research proposals on incorporating Privacy Policy Specification into the mainstream requirements and software engineering methodologies. It is attempting to consider policies and requirements as sets of wishes or desires that can be formalized and analyzed for consistency [85]. For instance, Moffett [86] recognized the benefit of regarding high-level policies as requirements and low level policies as their implementation. Accordingly, he proposed to integrate the specification of privacy policies into the requirements specification process.

Antón et al. [51] have analyzed the relationship between privacy policies and requirements. They noticed that policies and requirements are similar because they express desire, rather than fact. Additionally, both policies and requirements typically specify what must be done. On the other hand, privacy policies provide

a high-level description of enterprise practices, rather than system functionalities as do requirements. Accordingly, they have extended GBRAM by addressing the development of privacy policies during goal and scenario-driven activities. Similarly to what has been done in GBRAM for system requirements, they proposed heuristics in order to offer a methodological and systematic approach to identifying and formulating privacy policies and guaranteeing that system requirements are compliant with these policies.

Based on this work, Antón et al. [84] have structured the privacy policy domain with goal taxonomies. They classified privacy policies as either *privacy protection goals* or *privacy vulnerability goals*. Privacy protection goals are related to privacy principles and data subject rights and subdivided, along the guidelines given in FIPs [87], in five categories, namely notice/awareness, choice/consent, access/participation, integrity/security, and enforcement/redress. Privacy vulnerability goals concern threats and are classified in seven categories, namely monitoring, aggregation, storage, transfer, collection, personalization, and contact. These taxonomies has been integrated in the Evolutionary Prototyping with Risk Analysis and Mitigation (EPRAM) framework [88], which combines evolutionary prototyping with risk mitigation techniques for verifying the compliance of system requirements with security and privacy policies.

8.2. Aligning enterprise goals and data protection policies

Recent years have seen the emergency of dealing with security and privacy from the early phases of the system development process [31,80]. Unfortunately, security and privacy modeling has been largely independent of system requirements and system models. The usual approach towards the inclusion of security and privacy within a system is to identify security requirements after system design [89]. This is a critical problem, mainly because protection mechanisms have to be fitted into a pre-existing design which may not be able to accommodate them [90]. Consequently, security, privacy, and functional requirements can be misaligned or, even worst, in conflict with one another.

One of the current research challenges is to integrate security and privacy requirements analysis with the standard requirements engineering process. Jürjens [91] used stereotype {rbac} to specify the need of a protection mechanism based on RBAC [71] in the systems. Basin et al. [92] proposed a modeling language based on UML to model access control policies. Similarly, Doan et al. [93] proposed a metamodel to incorporate Mandatory Access Control (MAC) [68] in UML and Ray et al. [94] proposed to model RBAC policies as patterns in UML diagram template. However, these proposals do not provide any methodological support for driving system designers in the definition of access control and data protection policies.

Some researchers have already investigated the scope and applicability of access control policy languages and examined how organizational requirements can be captured by them. Neumann et al. [95] proposed a scenario-driven role engineering process for defining RBAC policies. This process starts by identifying usage scenarios where actions and events are seen as steps. These scenarios are successively used to derive the access operations, which are necessary to execute a sequence of steps, and related security constraints, such as separation of duty constraints. The identified access operations and constraints are stored in a permission catalog and constraints catalog, respectively. These catalogs are used to derive a preliminary role hierarchy that is at the basis of the concrete RBAC model. Liu et al. [80] showed how to specify an access control requirements model respecting certain security and privacy properties. To this intent, they proposed to rewrite i models in Alloy [96] and verify whether the defined access control policies satisfy least privilege and separation of duty properties. However, this approach is not able to capture the granularity de-

manded by an access control policies. For instance, a general Alloy function `need_access` is used to link a task with the permission on resources, which are necessary to execute a task. This function, however, does not allow one to distinguish the particular rights (e.g., read, write, etc.) that a user must have to execute the task.

Though some proposals for defining access control policies on the basis of the requirements model can be found in the literature, proposals addressing the definition of data protection policies are far less frequent. The only work we are aware is the one proposed by He et al. [52], who extended GBRAM by presenting a goal-driven framework for modeling privacy requirements in the role engineering [97]. This framework includes a context-based data model, in which privacy elements (i.e., purposes, conditions, obligations, retention period, and data recipient) are represented as attributes of roles, permissions, and objects, and a goal-driven role engineering process, which addressed how the privacy contexts in the data model can be elicited and modeled. Data protection requirements are thus modeled as contexts and constraints of permissions and roles. As the original GBRAM, this framework does not provide facilities for an accurate analysis of the structure of an organization and its environmental setting. Consequently, it is not able to capture, for instance, the different nature of obligations.

8.3. Aligning privacy policies, user preferences, and data protection policies

Organizations disclose their privacy promises by means of privacy policies. Unfortunately, this is not enough to safeguard data subject rights and privacy principles because privacy policies do not address the problem of how personal data are actually handled after collection. One of the challenges in managing privacy is thus to ensure that the promises made by enterprises to customers are actually enforced [98,99,26].

This problem has been partially addressed by Agrawal et al. [23] in the definition of their Hippocratic database systems. They provided a privacy metadata schema compounded by two tables, namely privacy-policies table and privacy-authorization table. The former captures the privacy policy by defining purpose, external recipient and retention period for each piece of data, whereas the latter describes the access control that supports the privacy policy by defining purpose and authorized users. Before a user provides any information, privacy policies are matched with user preferences. Data are thus inserted in the database only if the privacy policy do not violated user preferences. The purpose combined with the information in the privacy-authorization table is used to restrict access. LeFevre et al. [100] enhanced Hippocratic databases for enforcing queries to respect privacy policies and user preferences. In essence, they proposed to enforce the minimal disclosure principle by providing mechanisms that control who can access their personal data and for which purpose. However, these proposals do not guarantee the consistency between privacy policies and data protection policies.

Other researchers attempt to connect privacy policy languages, such as P3P, with Privacy-Aware Access Control languages, such as EPAL and XACML. We noticed that languages for Privacy-Aware Access Control are often coupled with mechanisms for translating data protection policies into privacy policies. For instance, Karjoth et al. [99] provided support for an automatic transformation of privacy practices expressed in E-P3P into privacy promises expressed in P3P. In this work, the authors recognized that E-P3P policies and, in general, data protection policies are finer-grained than what required by privacy policies. Thereby, they may result to be too complex for end-users. Accordingly, they proposed a transformation method from fine-grained E-P3P policies into coarse-grained privacy policy in P3P. Similarly, Barth et al. [98] defined an algorithm to translate DPAL policies into P3P policies. However, this approach

has a main drawback. Privacy policies should reflect the business strategies of an enterprise along with the entire range of alternatives that can be adopted by the enterprise itself to provide services. Conversely, data protection policies govern the access to data according to the particular alternative selected by a customer. Therefore, deriving privacy policies from data protection policies makes it not possible to capture all promises made by the enterprise.

Other approaches offer translation methods that work on other way round, that is, they define data protection from privacy policies and user preferences. Along this direction, Massacci et al. [25] proposed an approach for creating data protection policies from privacy policies and user preferences, ensuring Minimal Disclosure. This framework allows virtual organizations to model purpose hierarchies using weighted directed acyclic hypergraphs and specify the data items needed to satisfy the leaf purposes and customers to express their preferences in the form of privacy penalties associated with each personal data item and each partner participating to the virtual organization. Minimum weight traversal algorithms are then used to determine the process to deliver the desired service with the smallest privacy penalty. The calculated path is used to define the minimum set of authorizations necessary to achieve root purposes according to user preferences. However, this approach, as well as Hippocratic databases, requires privacy policy to have the same level of granularity demanded by data protection policies.

8.4. Privacy alignment and compliance discussion

As we have seen in previous sections, there are several policy description languages that have been used for privacy policy and data protection policy specifications as well as some requirements engineering methodologies addressing the analysis of privacy concerns. In particular, the work on policy specification (both privacy and data protection) seems to be enough mature. Indeed, most policy languages offer the right constructs to capture privacy protection aspects. What is missing is a methodological support that explains how organizational requirements can be captured by such languages [101]. In other words, the connection between above activities (i.e., requirements analysis and policy specification) is not well established.

We attribute this to the lack of requirements engineering methodologies able to deal with the entire privacy domain. The alignment of privacy artifacts can be established only using a requirements engineering methodology able to capture all the aspects demanded by privacy and data protection policies. There is evidence in the literature that Requirements Engineering can support the specification of privacy policies [51] as well as of data protection policies [52] and, in general, access control policies [80,95]. However, their drawbacks are mainly due to the underlying requirements engineering methodology.

This need also raises when aligning privacy and data protection policies. As we have seen in Section 8.3, there are intrinsic factors that make it impractical the direct transformation of privacy policies into data protection policies and/or vice versa. Perhaps the main reason is that privacy and data protection policies address different problems. Consequently, the information they provide is different. For instance, privacy policies should contain data subject rights in respect of the data processing, whereas data protection policies do not. Moreover, the same information can be expressed differently or require a different level of granularity. For instance, the retention period expressed as a data element in P3P is represented as an obligation in E-P3P [72].

We believe that Requirement Engineering can come to the rescue for a number of reasons. In particular, Requirements Engineering can offer a unifying view of socio-technical systems built using structured methodologies. Requirements, privacy policies, and data protection policies can be thus represented in this view at different levels of abstraction. Here, the different level of granularity de-

manded by privacy policies and data protection policies can be captured and assembled in a natural way. This unified view also allows system designers to check the consistency among privacy artifacts using the formal techniques offered by the chosen methodology.

9. Comparison with US regulations

A challenging question is how to guarantee data protection across countries that have a different foundation of privacy. In this section, we provide an overview of the US legal system on data protection, pointing out the main conceptual and practical differences with the EU legal system.

The more preminent difference is that privacy has not been recognized as a fundamental right in the US legal system. In such a legal system, privacy has been connected to some values that constitute and permeate common conscience, as individual freedoms, freedom of speak, pursuit of happiness, distinguishing between private and public life [102,103].

This has fueled a doctrinaire debate about the necessity to distinguish between the right to privacy and the concept of privacy. Among the rights protected by the US Constitution, the right to privacy is the most difficult to understand. The most obvious reason for such a difficulty is the fact that privacy is not explicitly dealt with by the US Constitution [103]. The concepts of privacy is characterized by a non-absoluteness and a lack of clear and unequivocal definition and contains several kinds of rights, as freedom of thought, right to be let alone, control on own data, freedom from wire-tapping, defense of reputation, etc. [104].

Theoreticians of privacy attempt to establish the common elements of privacy and, based on them, set up a debate about the conceptualization of privacy, suggesting a multitude of different theories and approaches ranging from the control of personal data flow and the capacity and freedom to make decisions on personal matters to a psychological aspect of privacy protection [105]. The protection of privacy is also conditioned by the epistemological approach used to “explain” the right to privacy. In summary, the doctrine has provided a number of approaches to regulate privacy protection. The more remarkable approaches are:

- *Property approach*, which proposes that the legal system protects privacy by establishing a property right on personal data. As consequence, this privacy model provides a person with the right to sell his own personal data [106]. Actually, there are dispositions on property right matters addressing privacy issues, such as the appropriation tort and trespass [107].
- *Contractual approach*, which proposes that the legal system facilitates and encourages the use of “privacy agreements” between two parties when there are no informative asymmetries and expensive bargaining costs [39,40]. Specific contractual provisions can be used to regulate the collection, processing, and diffusion of personal data. In certain contexts, courts have recognized as legitimate actions based on “breach of implied contract” or tort based on “implicit duties” once certain relationships are established. Actually, contracts often function as a way of sidestepping federal and state privacy laws. For instance, many organizations include the request of consent of their employees for drug testing as well as e-mail and workplace surveillance in employment contracts.
- *Approaches based on non-economic considerations*:
 - some approaches conceive privacy as a fundamental civil liberty interest and demand specific privacy regulations based on it [108];
 - others approaches focus on the cognitive limits of people to fully understand the risks connected to the diffusion of personal data to third parties and, consequently, assert that the legal system should provide some corrective measures [109].

Another main difference with the EU is the sectoral approach adopted by the US. The European privacy model is characterized by the identification of privacy as a fundamental right and is governed by a comprehensive regulation. Differently, the US privacy model is fragmented. Basically, the US relies on a mix of legislation, regulations, and self regulations that are founded on different sources of law, namely common law, constitutional law, statutory law, and international law. For instance, the First Amendment protects the right to speak anonymously as well as individuals from disclosing information about the groups to which they belong or contribute. The Fourth Amendment gives citizens the right “to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures” [110,111].

Starting from '70s, the US Congress has also enacted several statutes with the intent of protecting privacy in various (public and private) sectors of the society [112]. As consequences, the Federal Statutory Law results in a fragmented legislation regulating specific matters as federal government processing, schools, investigations, trial, telecommunications, health care, etc. The main federal statutes on privacy are presented in Table 3. These statutes establish different sets of privacy principles; each of these sets applies to a particular application domain. Table 4 summarizes how the EU privacy principles are covered by US federal statutes. However, a positive answer in such a table does not correspond to an exact match between the EU privacy principles and the ones defined in US federal statutes. Rather, we employ a weaker correspondence with the intended meaning that US federal statutes cope, in some way, with such principles. For the sake of space, we only present the privacy principles set by FIPs and their relationships with the EU principles. According to Smith [117], FIPs privacy principles can be summarized as follows:

- **Collection limitation:** an organization shall not collect personal data whose very existence is secret. This principle tackles some aspects of fair and lawful processing and minimality principles, as the collection of personal data should not intrude upon data subject's privacy and should be limited to the minimum necessary to satisfy the intended purpose.

Table 3
US Federal Statutes

The <i>Fair Credit Reporting Act</i> of 1970 (FCRA) [113] as amended by the <i>Fair and Accurate Credit Transactions Act</i> of 2003 (FACTA) [114], recognizes to the citizens a number of rights with respect to the processing and diffusion of personal data by credit institutions
The <i>Code of Fair Information Practices</i> of 1973 (FIPs) [87] provides a framework for privacy laws as well as the foundation of an organization's privacy policy – whether a private, public or not-for-profit organization
The <i>Privacy Act</i> of 1974 [115] provides to the citizens a number of rights with respect to federal databases. It is applied only to public agencies and bodies
The <i>Health Insurance Portability and Accountability Act</i> of 1996 (HIPAA) [116] represents the first complete regulations on medical data protection

Table 4
EU privacy principles in US federal statutes (Key: Yes – present, No – absent)

EU principles	FCRA	FIP	Privacy Act	HIPAA
Fair and Lawful Processing	Yes	Yes	Yes	Yes
Consent	Yes	Yes	Yes	Yes
Purpose Specification	Yes	Yes	Yes	Yes
Minimality	Yes	Yes	Yes	Yes
Minimal Disclosure	Yes	Yes	Yes	Yes
Information Quality	Yes	Yes	Yes	Yes
Data Subject Control	Yes	Yes	Yes	Yes
Sensitivity	Yes	No	Yes	Yes
Information Security	Yes	Yes	Yes	Yes

- **Disclosure:** an individual shall be able to determine which personal data are collected and how they are used. This principle is related to the consent principle and, in particular, to the data subject control principle, as the data subject shall be able to check and influence the processing of his personal data.
- **Secondary usage:** an individual shall be able to prevent that his personal data obtained for one purpose will be used or made available for other purposes without his explicit consent. This principle wraps purpose specification, minimal disclosure, and consent principles, as the data processing has to be linked to the specific purpose and the disclosure of data to third parties shall be carried on only when the data subject has given his consent.
- **Record correction:** an individual shall be able to correct or amend his personal data. This principle is related to the information quality principle, as personal data shall be accurate, relevant, and complete.
- **Security:** any organization creating, maintaining, using, or disseminating personal data shall assure the reliability of the data for their intended use and prevent their misuse. This principle corresponds to the information security principle, as organizations have to ensure an appropriate level of security when processing personal data.

Other differences between the EU and US legal systems can be found in the measures used to implement privacy principles and in the figures introduced by the legal system. The Principle of Consent, which is fundamental in the EU legal system, has been interpreted in a less restrictive way the US legal system. The EU has adopted an opt-in system and, in some particular cases, a more strict double opt-in system, whereas the US often adopts an opt-out system. Moreover, the US did not require the creation of government privacy authority agencies, registration of the data processing performed by an organization with those agencies, and, in the case of sensible data, a prior approval before their processing. As a result of these differences, the Directive could significantly affect the ability of US organization to engage in businesses with the EU. Indeed, the EU Directive forbids the transfer of personal data to a non-European Union country unless that country ensures an adequate level of privacy protection [3].

In order to close this gap, the US Department of Commerce in consultation with the European Commission developed the Safe Harbor [118]. This policy agreement provides US organizations with guidelines to simplify the procedure to comply with the EU Directive on data protection. Certifying to the Safe Harbor will assure US organizations to comply with the requirements and principles set by the EU Directive. The Safe Harbor is thus an important way for US organizations to avoid interruptions in their businesses with the EU or facing prosecution by European authorities.

It is worth noting that above considerations have to be reassessed in the light of September 11th, 2001. This event was caused by an organization secretly settled and efficaciously expanded in the fabric of American society, and obliged the legislator and the public opinion to reconsider the right to privacy of citizens. In particular, it has spurred US authorities to strengthen the entire machinery of security on a legal and operational level, heightening the contrast between public security and citizen privacy, especially, in digital data protection [119–121]. The result was the Usa Patriot Act¹¹ that sacrificed privacy in favor of national security.

¹¹ Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act – Pub. L. No. 107-56, 115 Stat. 272 (2001).

10. Discussion and final remarks

In the last years, it has been recognized the crucial role that legal requirements play in privacy-aware technologies. To put it somewhat provocatively, we claim that a system that does not meet legal requirements (established by privacy and data protection legislation) cannot be considered as a privacy-aware system. Therefore, the synergy between law and technology is the key to the successful development of privacy-aware systems. However, such a synergy is an ambitious objective, difficult to reach but absolutely necessary to pursue. A first challenge is to make legal experts and computer scientists to interact with each other. Another challenge is the identification of the concepts necessary to express privacy-related legal requirements.

This work is a first step toward filling the gap between law and technologies with the intent to support the design of privacy-aware systems. Firstly, we have identified the concepts necessary to capture and represent legal requirements established by the European data protection legislation. We have then reviewed the state of the art in Privacy Requirements Engineering, Privacy Policy Specification, and Privacy-Aware Access Control and the relationships among these research areas. This survey was intended to understand if existing proposals suffice to cope with the privacy issues raised by privacy legislation. This analysis has revealed that the work in Privacy Policy Specification and Privacy-Aware Access Control is quite mature and most proposals offer the right concepts to address privacy issues. On the contrary, Privacy Requirements Engineering is still immature. Though Requirements Engineering methodologies have the potentialities to assist policy writers in the specification of privacy and data protection policies, most proposals adopt “traditional” Requirements Engineering or Security Requirements Engineering frameworks and use such frameworks as they are to capture privacy requirements. The main problem is that those frameworks lack fundamental concepts specific to privacy. Thereby, they cannot be used to capture several situations that are frequent in the privacy domain but cumbersome if not impossible to express with existing constructs. Our claim is thus that Requirements Engineering methodologies can assist system designers in the development of privacy-aware systems, but they need to be enhanced at meta-level with constructs tailored to capture privacy-related legal requirements.

In summary, this work intends to serve as a technical reference for the development of Privacy Requirements Engineering methodologies aiming to ensure that deployed systems guarantee a sufficient level of privacy protection. In particular, we are interested in the development of methodologies that assists system designer in modeling and analyzing privacy concerns and legal requirements from the early phases of the system development process as well as policy writers in the specification of privacy and data protection policies by providing a means for deriving such policies from the requirements model.

Acknowledgments

This work has been partially funded by EU Commission through the SERENITY project and by the MIUR through the FIRB TOCALIT project.

References

- [1] S.D. Warren, L.D. Brandeis, The right to privacy, *Harvard Law Review* 4 (5) (1890) 193–220.
- [2] L.A. Bygrave, Data protection law: approaching its rationale, logic, and limits, *Information Law Series*, 10, Kluwer Law International, The Hague, 2002.
- [3] S. Room, Data Protection & Compliance in Context, BCS, 2007.
- [4] S. Kenny, J. Borking, The value of privacy engineering, *Journal of Information, Law and Technology* 4 (1) (2002).
- [5] A. Adams, The implications of users' multimedia privacy perceptions on communication and information privacy policies, in: *Proceedings of Telecommunications Policy Research Conference*, 1999.
- [6] J.R. Collmann, T. Cooper, Breaching the security of the Kaiser Permanente internet patient portal: the organizational foundations of information security, *Journal of the American Medical Informatics Association* 14 (2) (2007) 239–243.
- [7] E. Silverman, Loss of Protected Patient Information Real Danger for Health Care Plans, *Management Care*.
- [8] I. Araujo, Privacy mechanisms supporting the building of trust in e-commerce, in: *Proceedings of International Workshop on Privacy Data Management*, IEEE Press, 2005, p. 1193.
- [9] L. Cranor, M. Langheinrich, M. Marchiori, J. Reagle, The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Recommendation (April 2002). Available from: <<http://www.w3.org/TR/P3P/>>.
- [10] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu, An XPath-based preference language for P3P, in: *Proceedings of WWW'03*, ACM Press, 2003, pp. 629–639.
- [11] L. Cranor, M. Langheinrich, M. Marchiori, A P3P Preference Exchange Language 1.0 (APPEL1.0), W3C Working Draft, 2002. Available from: <<http://www.w3.org/TR/P3P-preferences/>>.
- [12] P. Ashley, S. Hada, G. Karjoth, C. Powers, M. Schunter, Enterprise Privacy Authorization Language (EPAL 1.1), Research Report 3485, IBM Research (2003). Available from: <http://www.zurich.ibm.com/security/enterprise-privacy/e_pal/>.
- [13] A. Barth, J.C. Mitchell, J. Rosenstein, Conflict and combination in privacy policy languages, in: *Proceedings of WPES'04*, ACM Press, 2004, pp. 45–46.
- [14] J.-W. Byun, E. Bertino, N. Li, Purpose based access control of complex data for privacy protection, in: *Proceedings of SACMAT'05*, ACM Press, 2005, pp. 102–110.
- [15] M. Hilty, D.A. Basin, A. Pretschner, On obligations, in: *Proceedings of ESORICS'05*, LNCS, vol. 3679, Springer-Verlag, 2005, pp. 98–117.
- [16] H.J. Smith, Privacy policies and practices: inside the organizational maze, *CACM* 36 (12) (1993) 104–122.
- [17] B. Nusejbeh, S. Easterbrook, Requirements engineering: a roadmap, in: *Proceedings of ICSE'00*, ACM Press, 2000, pp. 35–46.
- [18] J. Dumortier, C. Goemans, Legal challenges for privacy protection and identity management, *Security and Privacy in Advanced Networking Technologies*, vol. 193, IOS press, 2004.
- [19] R. Pardolesi (Ed.), *Diritto alla riservatezza e circolazione dei dati personali*, Giuffrè, 2003.
- [20] P. Carey, *Data Protection. A Practical Guide to UK and EU Law*, 2nd ed., Oxford University Press, 2004.
- [21] J.H. Saltzer, M.D. Schroeder, The protection of information in computer systems, *Proceedings of the IEEE* 63 (9) (1975) 1278–1308.
- [22] ISO/IEC, Code of practice for information security management, ISO/IEC 17799:2005, 2005.
- [23] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu, Hippocratic databases, in: *Proceedings of VLDB'02*, 2002, pp. 143–154.
- [24] E. Bertino, J.-W. Byun, N. Li, Privacy-preserving database systems, in: *FOSAD 2004/2005*, LNCS, vol. 3655, Springer-Verlag, 2005, pp. 178–206.
- [25] F. Massacci, J. Mylopoulos, N. Zannone, Hierarchical hippocratic databases with minimal disclosure for virtual organizations, *VLDBJ* 15 (4) (2006) 370–387.
- [26] G. Karjoth, M. Schunter, M. Waidner, Platform for enterprise privacy practices: privacy-enabled management of customer data, in: *Proceedings of PET'02*, LNCS, vol. 2482, Springer-Verlag, 2002, pp. 69–84.
- [27] J. Gordijn, M. Petit, R. Wieringa, Understanding business strategies of networked value constellations using goal- and value modeling, in: *Proceedings of RE'06*, 2006, pp. 126–135.
- [28] J. Mylopoulos, Aligning information strategy with business strategy – a technical perspective, in: *Proceedings of NGITS'93*, 1993.
- [29] F. Massacci, J. Mylopoulos, N. Zannone, From hippocratic databases to Secure Tropos: a computer-aided re-engineering approach, *IJSEKE* 17 (2) (2007) 265–284.
- [30] R. Clarke, e-Consent: a critical element of trust in e-Business, in: *Proceedings of the 15th Bled Electronic Commerce Conference*, 2002.
- [31] P. Giorgini, F. Massacci, N. Zannone, Security and trust requirements engineering, in: *FOSAD 2004/2005*, LNCS, vol. 3655, Springer-Verlag, 2005, pp. 237–272.
- [32] N. Nagaratnam, D. Lea, Practical delegation for secure distributed object environments, *Distributed Systems Engineering* 5 (4) (1998) 168–178.
- [33] H. Gomi, M. Hatakeyama, S. Hosono, S. Fujita, A delegation framework for federated identity management, in: *Proceedings of DIM'05*, ACM Press, 2005, pp. 94–103.
- [34] E. Bertino, P. Samarati, S. Jajodia, An extended authorization model for relational databases, *TKDE* 9 (1) (1997) 85–101.
- [35] B.S. Firozabadi, M.J. Sergot, Revocation schemes for delegated authorities, in: *Proceedings of POLICY'02*, IEEE Press, 2002, pp. 210–213.
- [36] D. Wijesekera, S. Jajodia, F. Parisi-Presicce, A. Hagström, Removing permissions in the flexible authorization framework, *TODS* 28 (3) (2003) 209–229.
- [37] L. Zhang, G.-J. Ahn, B.-T. Chu, A rule-based framework for role-based delegation and revocation, *TISSEC* 6 (3) (2003) 404–441.
- [38] P. Zheng, Tradeoffs in certificate revocation schemes, *ACM SIGCOMM Computer Communication Review* 33 (2) (2003) 103–112.

- [39] S.A. Bibas, A contractual approach to data privacy, *Harvard Journal of Law & Public Policy* 12 (2) (1994) 591–622.
- [40] S. Shorr, Personal information contracts: how to protect privacy without violating the first amendment, *Cornell Law Review* 80 (1995) 1756–1793.
- [41] N. Damianou, N. Dulay, E. Lupu, M. Sloman, The Ponder policy specification language, in: *Proceedings of POLICY'01*, LNCS, vol. 1995, Springer-Verlag, 2001, pp. 18–39.
- [42] C. Bettini, S. Jajodia, X.S. Wang, D. Wijesekera, Obligation monitoring in policy management, in: *Proceedings of POLICY'02*, IEEE Press, 2002, pp. 2–12.
- [43] C. Bettini, S. Jajodia, X.S. Wang, D. Wijesekera, Provisions and obligations in policy management and security applications, in: *Proceedings of VLDB'02*, 2002, pp. 502–513.
- [44] OASIS, eXtensible Access Control Markup Language (XACML) Version 2.0, OASIS Standard, 2005. Available from: <http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf/>.
- [45] R. Hilpinen, Deontic logic, in: *The Blackwell Guide to Philosophical Logic*, Blackwell Publishers Ltd., 2001, pp. 159–182.
- [46] J.R. Searle, *Speech Acts: An Essay in the Philosophy of Language*, Cambridge University Press, 1969.
- [47] J. Park, R. Sandhu, The UCONABC usage control model, *TISSEC* 7 (1) (2004) 128–174.
- [48] E. Bertino, C. Bettini, E. Ferrari, P. Samarati, A temporal access control mechanism for database systems, *TKDE* 8 (1) (1996) 67–80.
- [49] E. Bertino, C. Bettini, E. Ferrari, P. Samarati, An access control model supporting periodicity constraints and temporal reasoning, *TODS* 23 (3) (1998) 231–285.
- [50] E. Bertino, P.A. Bonatti, E. Ferrari, TRBAC: a temporal role-based access control model, *TISSEC* 4 (3) (2001) 191–233.
- [51] A.I. Antón, J.B. Earp, Strategies for developing policies and requirements for secure E-Commerce systems, in: *E-Commerce Security and Privacy*, Kluwer Academic Publishers, 2001, pp. 29–46.
- [52] Q. He, A.I. Antón, A framework for modeling privacy requirements in role engineering, in: *Proceedings of REFSQ'03*, 2003, pp. 137–146.
- [53] F. Massacci, N. Zannone, Privacy is linking permission to purpose, in: *Proceedings of the 12th International Workshop on Security Protocols*, vol. 3957, LNCS, Springer-Verlag, 2004, pp. 179–191.
- [54] J. Camenisch, A. Shelat, D. Sommer, S. Fischer-Hübner, M. Hansen, H. Krasemann, G. Lacoste, R. Leenes, J. Tseng, Privacy and identity management for everyone, in: *Proceedings of DIM'05*, ACM Press, 2005, pp. 20–27.
- [55] S. Dritsas, D. Gritzalis, C. Lambrinoukakis, Protecting privacy and anonymity in pervasive computing: trends and perspectives, *Telematics and Informatics* 23 (2).
- [56] U. Jendricke, M. Kreutzer, A. Zugenmaier, Pervasive privacy with identity management, in: *Proceedings of UbiComp'02*, 2002.
- [57] J. Feigenbaum, M.J. Freedman, T. Sander, A. Shostack, Privacy engineering for digital rights management systems, in: *Proceedings of DRM'01*, Springer-Verlag, 2002, pp. 76–105.
- [58] M. Backes, G. Karjoth, W. Bagga, M. Schunter, Efficient comparison of enterprise privacy policies, in: *Proceedings of SAC'04*, ACM Press, 2004, pp. 375–382.
- [59] US Federal Trade Commission, *Privacy Online: A Report to Congress* (1998). Available from: <<http://www.ftc.gov/reports/privacy3/>>.
- [60] Garante per la protezione dei dati personali, *Privacy Policy* (2003). Available from: <<http://www.garanteprivacy.it/garante/doc.jsp?ID=36573/>>.
- [61] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu, An implementation of P3P using database technology, in: *Proceedings of the 9th International Conference on Extending Database Technology*, LNCS, vol. 2992, Springer-Verlag, 2004, pp. 845–847.
- [62] G. Hogben, A technical analysis of problems with P3P 1.0 and possible solutions, in: *Proceedings of W3C Workshop on the Future of P3P*, 2002.
- [63] M. Schunter, E. Van Herreweghen, M. Waidner, Expressive privacy promises – how to improve the platform for privacy preferences (P3P), in: *Proceedings of W3C Workshop on the Future of P3P*, 2002.
- [64] T. Yu, N. Li, A.I. Antón, A formal semantics for P3P, in: *Proceedings of SWS'04*, ACM Press, 2004, pp. 1–8.
- [65] L.F. Cranor, J.R. Reidenberg, Can user agents accurately represent privacy notices?, in: *Proceedings of TPRC'02*, 2002.
- [66] N. Li, T. Yu, A.I. Antón, A semantics based approach to privacy languages, *Computer Systems Science and Engineering* 21 (5) (2006) 339–352.
- [67] A. Tumer, A. Dogac, H. Toroslu, A Semantic based privacy framework for web services, in: *Proceedings of ESSW'03*, 2003.
- [68] D.E. Denning, P.J. Denning, Certification of programs for secure information flow, *CACM* 20 (7) (1977) 504–513.
- [69] D. Downs, J. Rub, K. Kung, C. Jordan, Issues in discretionary access control, in: *Proceedings of Symposium on Security and Privacy*, IEEE Press, 1985, pp. 208–218.
- [70] S. Jajodia, P. Samarati, M.L. Sapino, V.S. Subrahmanian, Flexible support for multiple access control policies, *TODS* 26 (2) (2001) 214–260.
- [71] R.S. Sandhu, E.J. Coyne, H.L. Feinstein, C.E. Youman, Role-based access control models, *IEEE Computer* 29 (2) (1996) 38–47.
- [72] G. Karjoth, M. Schunter, A privacy policy model for enterprises, in: *Proceedings of CSFW'02*, IEEE Press, 2002, pp. 271–281.
- [73] M. Backes, B. Pfitzmann, M. Schunter, A toolkit for managing enterprise privacy policies, in: *Proceedings of ESORICS'03*, LNCS, vol. 2808, Springer-Verlag, 2003, pp. 162–180.
- [74] OASIS, Privacy policy profile of XACML v2.0, OASIS Standard (2005). Available from: <http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-privacy_profile-spec-os.pdf/>.
- [75] A.H. Anderson, A comparison of two privacy policy languages: EPAL and XACML, in: *Proceedings of SWS'06*, ACM Press, 2006, pp. 53–60.
- [76] P. Mazzoleni, B. Crispo, S. Sivasubramanian, E. Bertino, XACML policy integration algorithms, *TISSEC* 11 (1) (2008) 1–29.
- [77] P. Zave, Classification of research efforts in requirements engineering, *CSUR* 29 (4) (1997) 315–321.
- [78] E. Kavakli, C. Kalloniatis, P. Loucopoulos, S. Gritzalis, Incorporating privacy requirements into the system design process: the pris conceptual framework, *Internet Research* 16 (2).
- [79] L.K. Chung, B.A. Nixon, E.S.K. Yu, J. Mylopoulos, Non-Functional Requirements in Software Engineering, Kluwer Publishing, 2000.
- [80] L. Liu, E.S.K. Yu, J. Mylopoulos, Security and privacy requirements analysis within a social setting, in: *Proceedings of RE'03*, IEEE Press, 2003, pp. 151–161.
- [81] J. Mylopoulos, L. Chung, S. Liao, H. Wang, E. Yu, Exploring Alternatives During Requirements Analysis, *IEEE Software* 18 (1) (2001) 92–96.
- [82] H. Kaindl, A design process based on a model combining scenarios with goals and functions, *IEEE Transactions on Systems, Man, and Cybernetics* 30 (5) (2000) 537–551.
- [83] A.I. Antón, C. Potts, The use of goals to surface requirements for evolving systems, in: *Proceedings of ICSE'98*, IEEE Press, 1998, pp. 157–166.
- [84] A.I. Antón, J.B. Earp, A requirements taxonomy for reducing Web site privacy vulnerabilities, *REJ* 9 (3) (2004) 169–185.
- [85] A.I. Antón, J.B. Earp, R.A. Carter, Precluding incongruous behavior by aligning requirements with security and privacy policies, *Information & Software Technology* 45 (14) (2003) 967–977.
- [86] J. Moffett, Requirements and policies, in: *Proceedings of POLICY'99*, 1999.
- [87] US Department of Health, Education and Welfare, *The Code of Fair Information Practices*, 1973.
- [88] R.A. Carter, A.I. Antón, L. Williams, A. Dagnino, Evolving beyond requirements creep: a risk-based evolutionary prototyping model, in: *Proceedings of RE'01*, IEEE Press, 2001, pp. 94–101.
- [89] R. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, Wiley, 2001.
- [90] W. Stallings, *Cryptography and Network Security: Principles and Practice*, Prentice-Hall, Englewood Cliffs, NJ, 1999.
- [91] J. Jürjens, *Secure Systems Development with UML*, Springer-Verlag, 2004.
- [92] D. Basin, J. Doser, T. Lodderstedt, Model driven security: from UML models to access control infrastructures, *TOSEM* 15 (1) (2006) 39–91.
- [93] T. Doan, S. Demurjian, T.C. Ting, A. Ketterl, MAC and UML for secure software design, in: *Proceedings of FMSE'04*, ACM Press, 2004, pp. 75–85.
- [94] I. Ray, N. Li, R. France, D.-K. Kim, Using UML to visualize role-based access control constraints, in: *Proceedings of SACMAT'04*, ACM Press, 2004, pp. 115–124.
- [95] G. Neumann, M. Strembeck, A scenario-driven role engineering process for functional RBAC roles, in: *Proceedings of SACMAT'02*, ACM Press, 2002, pp. 33–42.
- [96] D. Jackson, Alloy: a lightweight object modelling notation, *TOSEM* 11 (2) (2002) 256–290.
- [97] E.J. Coyne, Role engineering, in: *Proceedings of RBAC'95*, ACM Press, 1995, pp. 15–16.
- [98] A. Barth, J.C. Mitchell, Enterprise privacy promises and enforcement, in: *Proceedings of the 2005 Workshop on Issues in the Theory of Security*, ACM Press, 2005, pp. 58–66.
- [99] G. Karjoth, M. Schunter, E.V. Herreweghen, Translating privacy practices into privacy promises – how to promise what you can keep, in: *Proceedings of POLICY'03*, IEEE Press, 2003, p. 135.
- [100] K. LeFevre, R. Agrawal, V. Ercegovac, R. Ramakrishnan, Y. Xu, D.J. DeWitt, Limiting disclosure in hippocratic databases, in: *Proceedings of VLDB'04*, Morgan Kaufmann, 2004, pp. 108–119.
- [101] R. Crook, D. Ince, L. Lin, B. Nuseibeh, Security requirements engineering: when anti-requirements hit the fan, in: *Proceedings of RE'02*, IEEE Press, 2002, pp. 203–205.
- [102] P.M. Regan, *Legislating Privacy: Technology, Social Values, and Public Policy*, University of North Carolina Press, 1995.
- [103] D.J. Solove, M. Rotenberg, P.M. Schwartz, *Information Privacy Law*, 2nd ed., Aspen Publishers, 2006.
- [104] H. Gross, *The Concept of Privacy*, New York University Law Review 42.
- [105] J.C. Inness, *Privacy, Intimacy and Isolation*, Oxford University Press, 1992.
- [106] P. Samuelson, Privacy as intellectual property?, *Stanford Law Review* 52 (5) (2000) 1125–1173.
- [107] R.S. Murphy, Property rights in personal information: an economic defense of privacy, *Georgetown Law Journal* 84 (1996) 2381–2417.
- [108] S.G. Davies, Re-engineering the right to privacy: how privacy has been transformed from a right to a commodity, in: *Technology and Privacy: The New Landscape*, MIT Press, 1997, pp. 143–145.
- [109] K.J. Strandburg, Privacy, rationality, and temptation: a theory of willpower norms, *Rutgers Law Review* 57 (4).
- [110] O. Kerr, The fourth amendment and new technologies: constitutional myths and the case for caution, *Michigan Law Review* 102 (2004) 801–888.
- [111] D.A. Sklansky, The fourth amendment and common law, *Columbia Law Review* 100 (2000) 1739.
- [112] M. Rotenberg (Ed.), *The Privacy Law Sourcebook*, EPIC, 2000.

- [113] US Federal Trade Commission, The Fair Credit Reporting Act, Pub. L. No. 90-32, 15 U.S.C. §1681 et seq., 1970.
- [114] US Congress, The Fair and Accurate Credit Transactions Act, Pub. L. 108-159, 2003.
- [115] US Department of Justice, Privacy Act (1974). Available from: <<http://www.usdoj.gov/foia/privstat.htm>>.
- [116] US Congress, Health Insurance Portability and Accountability Act, Pub. L. No. 104-191, 1996.
- [117] R.E. Smith, The law of privacy in a nutshell, *Privacy Journal* (1993) 50–51.
- [118] US Department of Commerce, Safe Harbor Privacy Principles, 2000. Available from: <<http://www.export.gov/safeharbor/SHPRINCIPLESFINAL.htm>>.
- [119] D. Carpenter, Keeping secrets, *Minnesota Law Review* 86 (6) (2002) 1097–1114.
- [120] M. Rotenberg, Privacy and secrecy after September 11, *Minnesota Law Review* 86 (6) (2002) 1115–1136.
- [121] D.J. Solove, Privacy and power: computer databases and metaphors for information privacy, *Stanford Law Review* 53 (2001) 1393–1462.