

Protection Goals for Privacy Engineering

Marit Hansen, Meiko Jensen, and Martin Rost

Unabhängiges Landeszentrum für Datenschutz Schleswig-Holstein (ULD)
Kiel, Germany

Email: Meiko.Jensen@rub.de, {Marit.Hansen | Martin.Rost}@datenschutzzentrum.de

Abstract—Six protection goals provide a common scheme for addressing the legal, technical, economic, and societal dimensions of privacy and data protection in complex IT systems. In this paper, each of these is analyzed for state of the art in implementation, existing techniques and technologies, and future research indications.

Keywords—privacy; data protection; protection goals; confidentiality; integrity; availability; unlinkability; transparency; intervenability; research challenges

I. INTRODUCTION

The first decade of the 21st century has seen a lot of technological progress that enabled governments and companies with a broad set of tools for collection, processing, and correlation of data. The amount of data produced in this society is about to increase rapidly, with most of it being generated by actions of human individuals. The rules how the data will—or at least can—be accessed and analyzed are defined by powerful organizations, e.g. by providers of social networks, search engines, or cloud computing, by infrastructure providers, and by governments including their secret services.

In this rush of technology, other essential societal elements did not keep pace. Most prominently, the legal and socio-economic dimensions of the digital society were evaluated slowly, if ever. In this context, the essential field of privacy and data protection has recently become one of the most severely damaged aspect of the digital society. For the Future Internet, this lack of societal compliance with the privacy and data protection needs is unacceptable, and must be addressed adequately.

In this aspect, recent research efforts have come up with a formalized model for incorporating privacy and data protection criteria in the design process: On the basis of six protection goals, engineers can derive requirements fitting for their use case, choose techniques and technologies to implement those requirements, and evaluate the privacy impacts and conditions of their IT systems. These protection goals (as proposed e.g. in [1], [2], [3], [4]) provide an interdisciplinary standard model for the development process as well as for assessing and judging the consequences of utilizing complex IT systems with respect to privacy and data protection.

This paper is about these six protection goals. Beyond the sheer listing and interrelation of the six protection goals themselves, the paper sheds a spotlight on state of the art techniques and technologies that can help implementing these six protection goals in smart services of the Future Internet. The intention of this paper is to give an overview and some pointers to ongoing research in this area, but it does not claim to be a complete list of techniques and technologies.

II. SIX PROTECTION GOALS FOR PRIVACY ENGINEERING

The original proposal of the six protection goals for privacy engineering distinguishes between three classic protection goals known for years within the IT security domain, and three protection goals genuine for privacy and data protection.

In this context, it is important to note that the terms of *privacy* and *data protection* are not synonyms. In the variety of definitions, cultural concepts, and translations, the clear distinction is blurred so that only few authors draw this strict line between the two terms. However, the EU Charter of Fundamental Rights clearly distinguishes between privacy and data protection (cf. [5]). This reflects a main difference: *privacy* usually takes the perspective of an individual who tries to fight back against the impertinence of control of others. *Data protection* rather refers to the organizational perspective, namely the social context of information processing, where self-determination and privacy are only possible if organizations are prevented from (mis-)using their power advantage over people. So, *data protection* tackles e.g. real choice in markets, functioning separation of powers in a constitutional state, democratic decision making, and free discourses. The data protection laws usually address this bigger objective by regulating the use of personal data, and thereby indirectly strengthening the fundamental rights in society.

The discipline of *privacy engineering* develops techniques and methods for both aspects: on the one hand, these techniques can be used for the domestication of organizations that deal with personal data, and on the other hand they provide immediately effective protection of the personal data of those concerned. Also, *information security* has to be recognized to support privacy engineering, but in this respect the often predominant focus on the interests of the organization (e.g. the service provider) has to be shifted towards the rights of the individual—very much as intended by the concept of multilateral security (cf. [6], [7]) that aims at empowering users, and stressing that imposing disadvantageous compromises on users must be prevented (cf. [8]).

Multilateral security research has highlighted not only the tensions between different stakeholders' interests in a system, but also influences of protection goals and subgoals on each other (cf. [9]). Indeed, the scheme of protection goals visualizes potential or factual conflicts, some of which we will depict later on.

In the following, we introduce protection goals from the fields of security as well as privacy and data protection.

A. Confidentiality, Integrity, and Availability

The traditional consideration of information security in IT systems has come up with a common set of three so-called *security protection goals* (see e.g. [10]), namely *confidentiality*, *integrity*, and *availability*. Known as the CIA triad, these three aspects are commonly considered of critical importance to evaluate an IT system’s security conditions. Therein, *confidentiality* addresses the need for secrecy, i.e. the non-disclosure of certain information to certain entities within the IT system in consideration. *Integrity* expresses the need for reliability and non-repudiation (cf. e.g. [11]) regarding a given piece of information, i.e. the need for processing unmodified, authentic, and correct data. As an important subset of such data, identity-related information is needed in authentic way to perform access control operations. *Availability* represents the need of data to be accessible, comprehensible, and processable in a timely fashion. Where confidentiality addresses non-disclosure to unauthorized entities, availability requires explicit and full disclosure to authorized entities—wherein the distinction between authorized and unauthorized entities typically needs integrity measures.

Each of these protection goals assumes an IT system beneath that supports or limits the particular protection goal, based on the technical details of its implementation. What is left unconsidered in this triad approach is the surrounding “real world”, i.e. the organizational and societal dimensions of the IT system, and especially its impact on the privacy of individuals (i.e. users, citizens, customers, patients, employees, administrators, etc.). Hence, for the purpose of creating a feature-complete model for evaluating an IT system’s impact on all aspects of privacy and data protection, these three security protection goals are complemented with three further protection goals on privacy and data protection that address the stated issues.

B. Unlinkability

The protection goal of *unlinkability* is defined as the property that privacy-relevant data cannot be linked across domains that are constituted by a common purpose and context (cf. [4]). This implies that processes have to be operated in such a way that the privacy-relevant data are not linkable to any privacy-relevant information outside of the domain.

Unlinkability is related to the requirements of necessity and data minimization as well as purpose determination, purpose separation, and purpose binding. The most effective method for unlinkability is data avoidance. Other methods for achieving or supporting unlinkability are e.g. data reduction, generalization, data hiding, separation, and isolation. The unlinkability protection goal should be considered already in early engineering phases because otherwise the design decisions taken (e.g. unique identifiers that enable context-spanning linkage, cf. [12]) may prevent a proper realization.

Note that this definition of unlinkability is much broader than in most terminology papers (cf. e.g. [13] for data minimization in communication systems) or in the Common Criteria Standard [14]. Those publications regard unlinkability as a specific property or goal of data minimization, with having other concepts such as anonymity or unobservability on the same level. Our broader definition addresses a generalized

view on which *item of interest* can be unlinkable to which other *item of interest* (cf. [13]). The definition above does not only tackle pieces of data, but also processes or domains that should not be linkable, so it abstracts from the scenario of communication systems. In this respect, the widened definition of unlinkability even encompasses societal concepts such as division of power. We believe that this feature is important for privacy engineering: It would not be sufficient to restrict the view on technical implementations only. It is necessary to consider the context and environment (including legal norms or societal values) for deriving requirements and for a proper realization in a world where mere technical solutions cannot provide satisfying answers.

C. Transparency

The protection goal of *transparency* is defined as the property that all privacy-relevant data processing—including the legal, technical, and organizational setting—can be understood and reconstructed at any time (cf. [4]). The information has to be available before, during, and after the processing takes place. Thus, transparency has to cover not only the actual processing, but also the planned processing (ex-ante transparency) and the time after the processing has taken place to know what exactly happened (ex-post transparency). The level of how much information to provide and how to communicate has to be adapted according to the capabilities of the target audience, e.g. the data-processing entity, the user, an auditor, or a supervisory authority.

Transparency is related to the requirement of openness (cf. [15]). Furthermore, it is a prerequisite for accountability. Standard methods for achieving or supporting transparency comprise logging and reporting, documentation of the data processing, or user notifications.

D. Intervenableity

The protection goal of *intervenability* is defined as the property that intervention is possible concerning all ongoing or planned privacy-relevant data processing (cf. [4]). In particular it applies to the individuals whose data are processed. The objective of intervenability consists of the effective enforcement of changes and corrective measures. As one example, intervenability reflects the individuals’ rights to rectification and erasure of data, the right to withdraw consent, and the right to lodge a claim or to raise a dispute to achieve remedy. Similar to the other protection goals, intervenability is relevant for other stakeholders, e.g. for data-processing entities and supervisory authorities to effectively influence or even stop the data processing. Think, e.g., of a cloud application where the personal data of a service’s customer has to be erased—here the service provider must be able to enforce this erasure in the cloud that is run by a third party.

Methods for achieving or supporting intervenability comprise implementation of dedicated services for intervention, definition of break-glass procedures, and means to override automated decisions.

E. The Three Axes

Having a closer look at the full set of protection goals, it can easily be seen that there is no possibility to ensure




Fig. 1. The six protection goals for privacy engineering.

100% of each of the goals simultaneously: If a system provides confidentiality, this implies that access to certain data is restricted for certain entities—thereby violating availability. Integrity conflicts intervenability, as the former disallows subsequent changes to the integrity-critical data and processes, and the latter requires exactly such ability for subsequent modifications. Transparency and unlinkability also turn out to be of conflicting nature, as the former intends to increase an understanding of the actual data processing, e.g. by logging the actions of users and administrators, and the latter tries to avoid such knowledge, as it may be misused for unintended linkage.

Each such conflict can typically be mitigated, depending on the particular IT system in consideration, but for the general model of the privacy and data protection protection goals, these three pairs of mutually affecting protection goals are represented as opponents. Hence, the full set of aspects is commonly represented as a star of three axes, each representing one pair of opposing protection goals (cf. Figure 1).

Beyond the three explicit conflict axes, there exist several other interrelations among these six protection goals. For instance, in order to effectively utilize the protection goal of intervenability, a basic understanding of the particular IT system’s functionality is required. Thus, it becomes necessary to obtain such information first—a transparency feature. Similarly, confidentiality and unlinkability have a lot in common. If information is not accessible, it can also not be linked in unintended ways.

On the conflicting side, transparency may conflict confidentiality, e.g. if disclosure of inner workings of an IT system may violate its secrecy assumptions (*security by obscurity*). Similarly, the protection goal of integrity in identity management scenarios requires reliable information on a dedicated entity’s identity (*access control*), but in the same instant may automatically allow for linkage of that identity to other contexts (*profiling*), which would harm the unlinkability protection goal.

As can be seen, additional conflicts and cooperations among different protection goals may arise, depending on the particular setting of the scenario in consideration. The protection goals give the advantage of making conflicts explicit, thereby urging the stakeholders involved in engineering to mitigate them by deciding on priorities within the protection goals depending on the use case and suitable mechanisms to realize them.

F. Implementations in Law

The six protection goals and the induced requirement of harmonization among these can be found in several stages of legal implementation all over Europe. For instance, the protection goal of transparency is directly implemented in the upcoming European General Data Protection Regulation, whereas the other protection goals can easily be derived from the regulation’s articles as well: the CIA protection goals are addressed by the demands for security, the protection goal of unlinkability is covered, among others, by purpose limitation and data minimization, and intervenability comprises data subject rights, data portability, and other control features such as consent. The full system of the six protection goals was implemented in the German federal state of Schleswig-Holstein as part of the state’s data protection act, and an effort of ISO standardization of this methodology is currently underway. Since the protection goals can be derived from law on the one hand, and at least security engineers are familiar with the concept of protection goals on the other hand, they facilitate bridging the legal and technical communities for gaining mutual understanding and cross-discipline collaboration, which is a basis for successful privacy engineering.

III. IMPLEMENTING THE PROTECTION GOALS FOR PRIVACY ENGINEERING

For the classic protection goals of the IT security domain, a common set of technologies, often of cryptographic nature, already exists. Despite some ongoing debates on how and which of these to deploy, their use is widely understood.

For the other three protection goals, however, the means of technical implementation are not that obvious, and to the best of our knowledge have not yet been analyzed scientifically. This lack is to be addressed further in this paper.

A. Confidentiality, Integrity, Availability

One technical realization of the confidentiality protection goal is commonly found in the domain of cryptography, i.e. the existing toolbox of encryption and decryption schemes. Most notably, symmetric schemes like AES or 3DES provide encryption and decryption using a secret key that is shared only by intended data recipients. Similarly, asymmetric encryption schemes like RSA and ElGamal provide secrecy based on a pair of a public and a secret key. Furthermore, access control enforcement contributes to the confidentiality protection goal as well.

As for confidentiality, the realization of integrity is prevalently based on a cryptographic scheme that allows for detection of modifications of data. Here, a list of cryptographic schemes for *digital signatures* exists, e.g. based on RSA or message authentication codes (MAC). Other means of realization of the integrity protection goal are based on redundancy and comparison, data verification, and—again—access control enforcement.

Unlike the previous two, availability cannot be realized by means of cryptography. In contrast, it can only be realized by adding redundancy to the system, e.g. by means of storing multiple copies of the same data in different storage locations. This way, if one of the copies is destroyed or altered, the other

copies remain intact, and the original data can be restored. Similarly, availability of services is enhanced by means of techniques like load balancing and virtualization. Ranging from simple local backups of data at rest to complex byzantine agreement protocols for redundant processes in clouds (cf. e.g. [16]), the set of tools here covers a broad range of parameters, each with their specific flaws and merits.

B. Unlinkability

Under the umbrella of this protection goal, a lot of commonly known properties are subsumed. For instance, unlinkability refers to the property of *anonymity* and its different aspects of realization. Therein, the human individual is allowed to use certain sorts of services without revealing her identity. This concept is quite close to that of pseudonymity, with the core difference that anonymous usage does not allow re-identification of a user at any stage. For pseudonymization, a (trusted) entity has the information about the link between a pseudonym and the related identity (e.g. by keeping a list or being able to use a transformation algorithm), so identity of an individual can be uncovered at a later stage—if allowed by that entity.

Both techniques have been implemented in multiple technologies, ranging from anonymization services for Internet users to anonymous messaging in telephone systems. Similarly, pseudonymization has found broad implementations in the realm of medical data processing. Also, it is commonly used in Internet services, e.g. by setting up multiple user accounts under fake names.

Special attention has to be paid to the feasibility of anonymization in real-world contexts. Recent studies (cf. [17], [18], [19]) have found that “anonymized” data sets still could be linked to the real identities of users, despite the lack of a dedicated list of identities. More advanced approaches, like that of *k-anonymity* (cf. [20]), provide better unlinkability, but are also way harder to implement reliably in real-world contexts.

Another typical technique for enhancing unlinkability consists in the implementation of access restrictions for data and processes. By means of dedicated access control enforcements, the set of authorized entities that may read or change personal data is limited. Thus, the overall threat of malicious linkage of such data to other sources of information is reduced to the scope of these authorized entities.

However, the idea of the unlinkability protection goal goes way beyond these technologies. One of the core challenges here is the external processing paradox: in order to use the services an external entity provides, it becomes necessary to provide the required input data for these services to that entity. Thus, that external entity learns the data in clear, and is able to maliciously keep and misuse that data afterwards. Not providing the required information, however, disqualifies the individual from using the offered services, which is also not satisfactory.

Here, manifold experimental and early cryptographical approaches have emerged during the last decades. Techniques of *secure computation* (cf. [21]) and *homomorphic encryption* (cf. [22]) support the processing of data without learning the data. For instance, it is possible to calculate the *larger-than*

relation $>$ for two values at an external entity without revealing the values themselves to that entity (cf. Yao’s millionaire’s problem, [23]). More advanced schemes like oblivious transfer, private information retrieval, and similar techniques exist for a broad range of other data processing tasks (see [24], [25], [26], [27], [22], [28], [29]), but each of these solutions still have teething troubles in terms of real-world usage. Some of the technologies show interdependencies with other protection goals, e.g. integrity: Both anonymity and authenticity needs can be realized by privacy-enhancing attribute-based credentials (cf. [30], [31]); or authenticity of a piece of information in the sense of a proven attribution to a person can be prevented when choosing plausible deniability encryption techniques (cf. [32]).

The protection goal of unlinkability also addresses information hiding mechanisms that aim at realizing unobservability and undetectability (cf. [13]), e.g. steganographic technologies where messages are unperceivedly embedded in other data (cf. [33]).

C. Transparency

The common techniques for fostering the protection goal of transparency are centered around *storing* and *delivering* information. For example, the former includes all sorts of documentation and logging techniques. Ranging from organization plans over system architecture handbooks to source code APIs, the set of documentation possibilities is quite broad, and each of these contributes to the protection goal of transparency. One essential characteristic here is that all aspects of a data-processing system are well-defined in advance, so that a system can never run into an undocumented state or perform an unforeseen action.

Beyond system documentation, dedicated and complete logging mechanisms play another essential role in the provisioning of effective transparency. For instance, whenever a customer interacts with a particular service of a complex IT system, that service’s logging component keeps track of those interaction events, and enqueues them in a dedicated logging subsystem.

The transparency aspect of information *delivery* includes transmission of the stored information to the relevant entities, such as to the affected human individuals themselves, but also to an organization’s management unit, a supervisory authority, or a dedicated auditor. This can be performed based on a dedicated request (*What do we store about person X?*), or proactively by sending notification messages (*A photo was uploaded to our servers by person Y*). Each of these mechanisms increases the amount of information a requesting entity gets about the complex IT system in consideration. Thereby, each of these contributes to the protection goal of transparency. Similarly, the establishment of a dedicated support service (e.g. a helpdesk, a ticket system, or a support line) fosters the degree of transparency within and from outside of an organization.

However, the state of the art in these technologies is not sufficient for providing the desirable degree of transparency for the Future Internet. For example, even though a web server component keeps track of incoming HTTP connections, these events are typically processed only by adding a single line in the web server’s access log file. Commonly, there is no linkage

to the identity of the user, nor is the context of the connection taken into consideration. Hence, when it comes to subsequent inquiries regarding an individual's activities, it is quite likely that the access log entry is not linked to the individual, and thus is skipped.

Unfortunately, there is no simple solution to this sort of lack, as do exist for the security protection goals. The use of cryptography does not contribute much to the protection goal of transparency. Even if so, it is mostly due to its property of integrity, or based on the availability of transparency-related information services.

The most promising technologies regarding transparency improvements are thus not based on the use of data modification techniques (like encryption or digital signatures), but require dedicated *transparency services* to be implemented alongside the core services of the particular IT systems (cf. [34]). These services then take care of storing, linking, aggregating, and providing all information required for achieving a sufficient level of transparency regarding the data of human individuals. Moreover, these services then need to be made accessible to the entitled entities in a usable and understandable way. As one example, they should be converted into a representation that is readable and suitable for any particular motivation a human user of a system might have regarding the inquiry of her personal information. Typically, no human being would like to wade through a huge pile of web server access log entries, just for trying to find that particular incident when something went wrong. Here, a more favorable approach would support the user in this task, e.g. by providing searching capabilities over the full set of related information available.

Moreover, reasonable techniques for data aggregation and representation have to be developed, in order to provide the user with relevant, easy-to-understand representations of the most important characteristics of an IT system (e.g. a graphical illustration of the flow of information, accompanied with the ability to filter the individual data fields that are processed therein). Also, standardized graphical user interfaces, e.g. based on icons (cf. [12], [35]) or standardized graphs can contribute a lot to the understanding and perception of an IT system's inner workings with respect to the data of a single human individual.

As it can be seen, even though a lot of basic transparency-enhancing technologies already exist (cf. [36], [37]), there still is room for improvement regarding the overall level of transparency an IT system can provide.

D. Intervenability

Unlike transparency, the protection goal of intervenability has way less technologies and techniques elaborated to the degree of daily use. The idea of intervenability is to enable direct action by entitled entities, such as the data-processing organization itself, a supervisory authority, or the affected human individual whose personal data is processed. For instance, this covers means to interdict data transmission to third parties, to correct errors within the data (e.g. fixing typos in phone numbers), to delete certain part of the data at any stage of processing, or to stop the processing altogether.

Typically, implementing intervenability is challenging, as the IT system in consideration has to be robust enough to cope with partial unavailability of data, with non-execution of parts of its processes, and with temporal delays caused by interfering actions of users. Hence, the workload of implementing a fully intervenability-friendly IT system often is a multitude of the workload of implementing the core system's functionality—inducing a multitude of implementation costs as well. This has led to companies neglecting the intervenability options for their customers. What typically can be found in terms of intervenability is a configuration menu for the user's core personal data, and a clerk-operated help desk. Interference with ongoing processes, however, is rarely implemented.

Here, a substantial amount of research challenges for the Future Internet can be identified. How can complex business processes cope with changes of the data in transit? How do they handle incomplete or missing data, or temporally caused inconsistencies of the data? How can customers become entitled to use their right to intervene, while reducing both costs and implementation workload for the companies?

Although the technological capabilities of modern IT systems can be improved by far, it requires appropriate incentives for their development and application. Legal obligation to provide a sufficient amount of intervenability would be such an incentive, and can already be perceived in contexts where the six protection goals for privacy engineering have become legal norms. Here, the nearby future provides a large playground for existential research challenges, paving the way for a standardized implementation of intervenability technologies, systems, and processes.

IV. DISCUSSION

In this section we iterate over some of the most important existing approaches to privacy engineering, and we discuss their interrelation to the six protection goals. Moreover, we provide an extensive list of real-world application examples.

A. Interrelation to other Approaches for Privacy Engineering

The six protection goals can perfectly be combined with the well-known Privacy by Design concept promoted by Ann Cavoukian (cf. [38]) and the International Conference of Data Protection and Privacy Commissioners (cf. [39]). In fact, the six goals, together with a catalog of techniques and technologies, provide the means to implement Privacy by Design appropriately and offer helpful guidance for engineers, developers, and evaluators.

Early privacy engineering approaches focused on confidentiality features (privacy should be preserved by preventing access to personal data) and data minimization (privacy should be preserved by not collecting personal data in the first place, or by deleting it as soon as possible). In this respect, the PriS method (cf. [40]) can be used as a basis for formal methods, or the LINDDUN privacy threat analysis framework (cf. [41]) may be employed to identify threats concerning so-called *hard privacy* (the field of the protection goals of confidentiality / unlinkability) and *soft privacy* for the user's content awareness as well as policy and consent compliance—the latter criteria partially refer to transparency and intervenability. The data minimizing approaches are still valid and have the advantage of

easier formalization, but fall short whenever data disclosure is inevitable. This was already noticed in the approach of Multi-lateral Privacy Requirements Analysis (cf. [42]) that addresses privacy as a control paradigm. It distinguishes between three types of privacy goals: *confidentiality goals*, *control goals*, and *practice goals*. These types of goals look similar to the *confidentiality / unlinkability*, *intervenability*, and *transparency* protection goals, but in their descriptions they do not cover legal demands such as the principle of purpose limitation or the data subject rights.

For the concept development and analysis phases of software development, eight privacy design strategies (cf. [43]) have been proposed. The four data-oriented strategies of *minimize*, *hide*, *separate* and *aggregate* are clearly related to the unlinkability and confidentiality protection goals. Of the four process-oriented strategies, *inform* and *demonstrate* reflect the transparency protection goal, *control* stands for intervenability. Finally, *enforce* is demanded by the integrity and availability protection goals. The privacy design strategies refine essential properties of the six protection goals, with a strong focus on the unlinkability protection goal, because that one is the first line of defense against misuse of data. Also, it reflects the maturity of the field of unlinkability (in particular data minimization, access control, and encryption methods), while techniques for transparency or intervenability are less standardized, less available, and often require—or are influenced by—non-technical system components, such as legal regulations, or organizational procedures. All in all, the privacy design strategies fit into the landscape of protection goals for privacy and data protection.

A first draft for the operationalization of privacy and data protection requirements using the six protection goals was created in 2012 with the *standardized data protection model* (SDM, cf. [44]). The SDM consists of three components: the six protection goals, the components of operations (data, IT, technical and organizational processes), and the protection requirements for the protection needs (categorized as *normal*, *high*, or *very high*). The model is based on the methodology of the IT Baseline Protection from the German Federal Office for Information Security (BSI, cf. [45]), which implements the ISO 27000 international standard series. However, the SDM model defines the need for protection from the perspective of the affected individuals and not—that is the main difference—from the business process point of view.

The process flow is as follows: First, the legal requirements must be clarified. In the European Union, for instance, the processing of personal data is by default prohibited by data protection law (EU Data Protection Directive 95/46/EC); this prohibition may be waived by e.g. a special law or a valid informed consent of the affected individual. If lawfulness of the data processing is given, in a second step the data are categorized according to the protection needs and, while considering the interests of all stakeholders, the protection requirements for each type of data are determined. These protection requirements then ascertain the required minimum protection measures to be applied when collecting and processing such data (cf. [46]). They also define the necessary security features of involved IT systems, and the technical and organizational processes they are embedded in.

B. Application Examples

Ambient Assisted Living (AAL) refers to technologies and services being provided with the aim of continuing to lead an independent life in one's own home (cf. [47]). Many of the proposed products and services use a comprehensive surveillance system by implementing cameras and sensors that can trigger an alarm at a guard service in case of an incident. In one study, the six protection goals were applied to seven such scenarios (cf. [48]). One interesting result was that several developers of AAL systems did not plan for intervenability: If a person living in an AAL-equipped home wanted to deactivate the monitoring video and sensors, e.g. because of visitors, this usually should be possible. However, since the guard services could not be held liable if they are not alerted due to a deactivated system, turning off the sensors would cause a shift in liability. Of course, this liability shift has to be clearly communicated to the inhabitants—as part of the transparency requirements, which are a challenge in any AAL system.

Similarly, *Cyber-Physical Systems (CPS)* demand solutions in particular for supporting unlinkability, transparency, and intervenability. In these systems, not all kinds of data are clearly personally identifiable, but the gathered information may affect the people's lives to a great extent. European data protection law does not fit well when it comes to CPS. Still, the application of the six protection goals (cf. [49]) yields viable results when deriving requirements for privacy engineering in cyber-physical systems.

Smart power grids and smart meters are regarded as enablers for environment-friendly and efficient provision and use of energy. Since personal data are being processed, privacy and data protection requirements have to be taken into account. The German Data Protection Commissioners on the Federal and the Land level have issued a guideline for data-protection compliant smart metering (cf. [50]). This guideline describes a method for determining the protection needs of different smart meter use cases and for applying the six protection goals in this context.

The German National IT Planning Council (“IT-Planungsrat”) is responsible for supporting interoperability in Germany's public administration. Several of the IT standards that are coordinated by this council comprise privacy and data protection issues. Therefore, some of its working groups started to use the six protection goals for defining the requirements for data exchange in e-government applications (cf. [51], [52]).

The lifecycle of *eID systems* was the focus of a research task performed in the European project “ABC4Trust – Attribute-based Credentials for Trust” (cf. [53]). Attribute-based credentials—other than traditional eID systems—can provide unlinkability. The options concerning unlinkability and the other protection goals were elaborated in each of the phases of the lifecycle of the eID systems. One interesting finding concerns debugging: while unlinkability is not necessarily helpful in an early test phase where errors have to be debugged, this should change when the system becomes more mature. But if a system design guaranteed unlinkability, testing of the functionality, but also of the privacy properties, is not trivial.

A thesis on privacy and security risk analysis of identity management systems took the six protection goals for elabo-

rating a set of risk factors (cf. [54]). The author points out that identifying and understanding conflicts between the protection goals are necessary for developing an adequate and a balanced risk analysis model.

In the European project “TClouds – Trustworthy Clouds”, different cloud-based components for improving privacy and resilience for critical infrastructures were developed. For evaluating the results, a *Data Protection Impact Assessment* was performed that was based on the six protection goals (cf. [55]). The methodology proved to be flexible according to the varying demands of the different use cases (hospital cloud with medical data, smart meter cloud). Also, the protection goals were driving the recommendations for technical and organizational measures of data protection and data security in cloud computing of the Art. 29 Data Protection Working Party that is composed of representatives of data protection authorities in each EU country and of the European Commission (cf. [56]).

The independent data protection authorities in Germany have founded a joint working group (*UAGSDM*) that is currently developing a catalog of agreed reference measures implementing the six protection goals. Much of the work in 2014 was devoted to the development of mapping data protection laws and the protection goals. The adoption of this catalog is scheduled for autumn 2015.

The six protection goals have evolved and gained maturity over the last five years when different groups from research, standardization, or supervision have been applying them for their respective needs: for highly regulated and well established contexts as well as for emerging technologies where the legal demands have not been spelled out. In most cases, interdisciplinary teams used the protection goals because they could be understood by all stakeholders. The protection goals can give guidance for elaborating criteria and choosing mechanisms, but only this paper maps a large variety of techniques to the goals and focuses on the privacy engineering community.

V. CONCLUSIONS AND RESEARCH INDICATIONS

The field of privacy and data protection in complex IT systems, e.g. in services of the Future Internet, provides a lot of highly prestigious research challenges for the decades to come. Being in the intersection of law, technology, computer science, and economics, the upcoming challenges in this field must be addressed sufficiently for each of these scientific areas.

In this interplay that requires cross-disciplinary comprehension among the stakeholders and engineers, the six protection goals have emerged as one of the most favorable schemes for measuring, assessing, implementing, and enforcing privacy and data protection in complex IT systems. We believe that we will see a rising attention in each of the listed sciences in the nearby future.

In this paper, a brief overview of the core challenges of each of these protection goals was given, accompanied with a glimpse at the state of the art in today’s real world systems. Based on the techniques outlined here, the implementation of Future Internet technologies for privacy and data protection can be envisioned, and the accompanying research challenges can be addressed.

ACKNOWLEDGMENT

We dedicate this work to the deceased Prof. Andreas Pfitzmann. As a committed and outspoken expert for privacy technologies, he laid the foundation for the six protection goals for privacy engineering. We thank him for his inspiring contributions and his constant willingness to improve the results in thorough discussions.

This work was partially funded by the European Commission, FP7 ICT program, under contract no. 318424 (FutureID project), and by the German Ministry of Education and Research within the Privacy-Forum.

REFERENCES

- [1] M. Rost and A. Pfitzmann, “Datenschutz-Schutzziele – revisited,” *Datenschutz und Datensicherheit*, vol. 33, no. 6, pp. 353–358, 2009.
- [2] M. Rost and K. Bock, “Privacy by Design and the New Protection Goals,” *EuroPriSe Whitepaper*, 2011. [Online]. Available: <https://www.european-privacy-seal.eu/results/articles/BockRost-PbD-DPG-en.pdf>
- [3] M. Hansen, “Top 10 Mistakes in System Design from a Privacy Perspective and Privacy Protection Goals,” in *Privacy and Identity for Life*, ser. IFIP AICT, vol. 375, IFIP International Federation for Information Processing. Springer, 2012, pp. 14–31.
- [4] G. Danezis, J. Domingo-Ferrer, M. Hansen, J.-H. Hoepman, D. Le Métayer, R. Tirtea, and S. Schiffner, “Privacy and Data Protection by Design – from policy to engineering,” ENISA, Tech. Rep., 2014. [Online]. Available: http://www.enisa.europa.eu/activities/identity-and-trust/library/deliverables/privacy-and-data-protection-by-design/at_download/fullReport
- [5] J. Kokott and C. Sobotta, “The distinction between privacy and data protection in the jurisprudence of the CJEU and the ECtHR,” *International Data Privacy Law*, vol. 3, no. 4, pp. 222–228, 2013.
- [6] K. Rannenberg, “Recent development in information technology security evaluation – the need for evaluation criteria for multilateral security,” in *Security and Control of Information Technology in Society – Proceedings of the IFIP TC9/WG 9.6 Working Conference*. North-Holland Publishers, 1994.
- [7] A. Pfitzmann, “Multilateral Security: Enabling Technologies and Their Evaluation,” in *Informatics – 10 Years Back, 10 Years Ahead*, ser. LNCS, R. Wilhelm, Ed., vol. 2000. Springer, 2001, pp. 50–62.
- [8] H. Bäumler, “IT-Sicherheit – für wen?” in *Mit Sicherheit in die Informationsgesellschaft – Tagungsband 5. Deutscher IT-Sicherheitskongress des BSI*, BSI, Ed. SecuMedia Verlag, 1997, pp. 481–490.
- [9] G. Wolf and A. Pfitzmann, “Properties of protection goals and their integration into a user interface,” *Comput. Netw.*, vol. 32, no. 6, pp. 685–699, May 2000. [Online]. Available: [http://dx.doi.org/10.1016/S1389-1286\(00\)00029-3](http://dx.doi.org/10.1016/S1389-1286(00)00029-3)
- [10] M. Bishop, *Introduction to Computer Security*. Addison-Wesley Professional, 2004.
- [11] A. McCullagh and W. Caelli. (2000, August) Non-Repudiation in the Digital Environment. First Monday, vol. 5. [Online]. Available: <http://firstmonday.org/ojs/index.php/fm/article/view/778>
- [12] M. Hansen, A. Schwartz, and A. Cooper, “Privacy and identity management,” *IEEE Security & Privacy*, vol. 6, no. 2, pp. 38–45, 2008.
- [13] A. Pfitzmann and M. Hansen, “A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management,” http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf, Aug. 2010, v0.34. [Online]. Available: http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf
- [14] *Common Criteria for Information Technology Security Evaluation – Part 2: Security functional components, Version 3.1, Revision 4*, CCMB Std., 2012. [Online]. Available: <http://www.commoncriteriaportal.org/files/ccfiles/CCPART2V3.1R4.pdf>
- [15] P. Birkinshaw, “Freedom of Information and Openness: Fundamental Human Rights?” *Administrative Law Review*, vol. 58, no. 1, pp. 177–218, 2006. [Online]. Available: <http://www.jstor.org/stable/40712007>

- [16] A. N. Bessani, M. Correia, B. Quaresma, F. André, and P. Sousa, "DepSky: Dependable and Secure Storage in a Cloud-of-Clouds," *TOS*, vol. 9, no. 4, p. 12, 2013.
- [17] M. Barbaro and T. Zeller, "A face is exposed for aol searcher no. 4417749," *The New York Times*, 2006. [Online]. Available: <http://select.nytimes.com/gst/abstract.html?res=F10612FC345B0C7A8CDDA10894DE404482>
- [18] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets," in *Security and Privacy, 2008. SP 2008. IEEE Symposium on*, May 2008, pp. 111–125.
- [19] Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, "Unique in the crowd: The privacy bounds of human mobility," in *Sci. Rep.* 3,1376, 2013, pp. 1–5.
- [20] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 5, pp. 557–570, 2002. [Online]. Available: <http://dx.doi.org/10.1142/S0218488502001648>
- [21] S. Micali and P. Rogaway, "Secure computation," in *Advances in Cryptology (CRYPTO 91)*, ser. Lecture Notes in Computer Science, J. Feigenbaum, Ed. Springer, 1992, vol. 576, pp. 392–404.
- [22] C. Gentry, "A fully homomorphic encryption scheme," Ph.D. dissertation, Stanford University, 2009.
- [23] A. C.-C. Yao, "Protocols for secure computations," in *FOCS*, vol. 82, 1982, pp. 160–164.
- [24] M. O. Rabin, "How to exchange secrets with oblivious transfer," 2005, Harvard University Technical Report 81. [Online]. Available: <http://eprint.iacr.org/2005/187>
- [25] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," *J. ACM*, vol. 45, no. 6, pp. 965–981, Nov. 1998. [Online]. Available: <http://doi.acm.org/10.1145/293347.293350>
- [26] F. Kerschbaum, A. Schröpfer, A. Zilli, R. Pibernik, O. Catrina, S. de Hoogh, B. Schoenmakers, S. Cimato, and E. Damiani, "Secure collaborative supply-chain management," *IEEE Computer*, vol. 44, no. 9, pp. 38–43, 2011.
- [27] M. Jensen and F. Kerschbaum, "Towards privacy-preserving XML transformation," in *ICWS*, 2011, pp. 65–72.
- [28] J. Camenisch, G. Neven, and M. Rückert, "Fully anonymous attribute tokens from lattices," in *SCN*, 2012, pp. 57–75.
- [29] J. Camenisch, "Cryptographic primitives for building secure and privacy respecting protocols," in *ACM Conference on Computer and Communications Security*, 2011, pp. 361–362.
- [30] J. Camenisch and A. Lysyanskaya, "An efficient system for non-transferable anonymous credentials with optional anonymity revocation," in *EUROCRYPT*, 2001, pp. 93–118.
- [31] S. Brands, *Rethinking Public Key Infrastructures and Digital Certificates; Building in Privacy*, 1st ed. MIT Press, 2000, ISBN 0-262-02491-8.
- [32] R. Canetti, C. Dwork, M. Naor, and R. Ostrovsky, "Deniable encryption," *Cryptology ePrint Archive*, Report 1996/002, 1996.
- [33] F. A. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Information Hiding – A Survey," in *Proceedings of the IEEE (special issue on protection of multimedia content)*, vol. 87, no. 7. IEEE, July 1999, pp. 1062–1078.
- [34] M. Jensen, "Towards privacy-friendly transparency services in inter-organizational business processes," in *Computer Software and Applications Conference Workshops (COMPSACW)*, 2013, pp. 200–205.
- [35] L.-E. Holtz, K. Nocun, and M. Hansen, "Towards displaying privacy information with icons," in *Privacy and Identity Management for Life*, ser. IFIP Advances in Information and Communication Technology, vol. 352. Springer, 2011, pp. 338–348.
- [36] T. Pulls, "Privacy-preserving transparency-enhancing tools," Ph.D. dissertation, Karlstad University, Department of Computer Science, 2012.
- [37] H. Hedbom, T. Pulls, and M. Hansen, "Transparency tools," in *Privacy and Identity Management for Life*, J. Camenisch, S. Fischer-Hübner, and K. Rannenberg, Eds. Springer, 2011, pp. 135–143.
- [38] A. Cavoukian. (2012, December) Operationalizing Privacy by Design: A Guide to Implementing Strong Privacy Practices. [Online]. Available: <http://www.privacybydesign.ca/content/uploads/2013/01/operationalizing-pbd-guide.pdf>
- [39] 32nd International Conference of Data Protection and Privacy Commissioners. (2010, October) Privacy by Design Resolution. 27-29 October 2010, Jerusalem, Israel. [Online]. Available: http://www.ipc.on.ca/site_documents/pbd-resolution.pdf
- [40] C. Kalloniatis, E. Kavakli, and S. Gritzalis, "Addressing Privacy Requirements in System Design: The PriS Method," *Requir. Eng.*, vol. 13, no. 3, pp. 241–255, Aug. 2008.
- [41] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, "A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements," *Requir. Eng.*, vol. 16, no. 1, pp. 3–32, Mar. 2011.
- [42] S. F. Gürses, "Multilateral privacy requirements analysis in online social networks," Ph.D. dissertation, Department of Computer Science, KU Leuven, Belgium, May 2010.
- [43] J. Hoepman, "Privacy Design Strategies – (extended abstract)," in *ICT Systems Security and Privacy Protection – 29th IFIP TC 11 International Conference, SEC*, 2014, pp. 446–459.
- [44] M. Rost, "Standardisierte Datenschutzmodellierung," *DuD – Datenschutz und Datensicherheit*, vol. 36, no. 6, pp. 433–438, 2012.
- [45] Bundesamt für Sicherheit in der Informationstechnik (German Federal Office for Information Security), "BSI Standard 100-2: IT-Grundschutz Methodology Version 2.0," *BSI Standards*, 2008. [Online]. Available: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/BSIStandards/standard_100-2_e.pdf
- [46] T. Probst, "Generische Schutzmaßnahmen für Datenschutz-Schutzziele," *DuD – Datenschutz und Datensicherheit*, vol. 36, no. 6, pp. 439–444, 2012.
- [47] DKE German Commission for Electrical, Electronic & Information Technologies of DIN and VDE. (2012, January) The German AAL Standardization Roadmap (= Ambient Assisted Living). [Online]. Available: <http://www.vde.com/en/dke/std/Documents/German%20AAL%20Standardization%20Roadmap%20%5B1%5D.pdf>
- [48] ULD, "Juristische Fragen im Bereich Altersgerechter Assistenzsysteme," ULD, Tech. Rep., 2011. [Online]. Available: <https://www.datenschutzzentrum.de/projekte/aal/>
- [49] E. Geisberger and M. Broy, Eds., *agendaCPS – Integrierte Forschungsagenda Cyber-Physical Systems*. Springer, 2012.
- [50] Conference of the Data Protection Commissioners of the Federal Government and the Länder and Düsseldorf Circle. (2012, June) Guideline for data-protection compliant smart metering. [Online]. Available: <http://www.bfdi.bund.de/SharedDocs/Publikationen/EN/NationaleDSK/GuidelineSmartMeter.pdf>
- [51] XKfz-Working Group, "Deutschland-Online Kfz-Wesen – Teilprojekt XKfz Standardisierung – XKfz-Spezifikation, Version 1.0.0," Ministerium des Innern, für Sport und Infrastruktur Rheinland-Pfalz, Tech. Rep., 2012. [Online]. Available: http://www.it-planungsrat.de/SharedDocs/Downloads/DE/Projekte/Abgeschlossene_Projekte/Kfz-Wesen/A6_Standardisierung_XKfz_Spezifikation.pdf?__blob=publicationFile
- [52] IT-Planungsrat, "Datenübertragung mit XTA (Version 2.0)," IT-Planungsrat, Tech. Rep., 2013. [Online]. Available: http://www.xoev.de/sixcms/media.php/13/Anlage2_XTA_Spezifikation.6863.pdf
- [53] H. Zwingelberg and M. Hansen, "Privacy Protection Goals and Their Implications for eID Systems," in *Privacy and Identity Management for Life*, ser. IFIP Advances in Information and Communication Technology, vol. 375. Springer, 2012, pp. 245–260.
- [54] E. Painsil, "Privacy and security risks analysis of identity management systems," Ph.D. dissertation, Gjøvik University College, 2013. [Online]. Available: http://brage.bibsys.no/xmlui/bitstream/id/101881/Thesis_Electronic_Version_EbenezerPainsil.pdf
- [55] N. Marnau, M. Jensen, E. Schlehahn, R. Morte Ferrer, and M. Hansen. (2013, October) Cloud Computing – Data Protection Impact Assessment. Deliverable D1.2.4 of the TClouds Project. [Online]. Available: <http://tclouds-project.eu/downloads/deliverables/TC-D1.2.4-Cloud-Computing-Privacy-Impact-Assessment-V1.1-Public.pdf>
- [56] Article 29 Data Protection Working Party, *Opinion 05/2012 on Cloud Computing*, 01037/12/EN, Article 29 Data Protection Working Party Std. WP 196, Adopted July 1st 2012. [Online]. Available: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp196_en.pdf