

Design for the Value of Privacy

Martijn Warnier, Francien Dechesne & Frances Brazier

Delft University of Technology

Abstract

In a time where more and more information about people is collected, especially in the digital domain, the right to be left alone and to be free of surveillance, i.e. privacy, is no longer as self-evident as it once was. Therefore it is important that new systems are designed with privacy in mind. This chapter explores the notion of privacy and how to design 'privacy preserving' systems: systems that are designed with privacy for the end-users in mind. Several design approaches that address this issue, such as 'Privacy by Design', 'Value Sensitive Design' and 'Privacy Enhancing Technologies' are discussed. Examples of privacy preserving (and breaking) systems, ranging from Smart Meters to Electronic Health records are used to illustrate the main difficulties of designing such systems.

Keywords

Privacy, Design, Value Sensitive Design, Smart Grid

1. Introduction

Throughout history only a privileged few enjoyed the privacy that in recent times has become more commonplace: the right to be left alone and not be under surveillance, both from peers as well as governments. In the last decades this has changed again with the rise of the Internet. What began as a means to freely and anonymously communicate with others around the world has become an instrument for violating the privacy of individuals at a scale hitherto not thought to be possible. Developments in information technology, such as increasing computing power, storage and communication, have led to many benefits for people, but individual privacy has come under threat. All kinds of data, ranging from marketing information (buyer profiling) to medical data are collected, linked and processed both by companies and governments. The increasing connectedness of stored data, makes it possible to link more data to individuals, thereby stretching what counts as 'personal data'.

The right to privacy (Warren and Brandeis, 1890) is a universal human right (Movios and Krup, 2009). It entails both freedom of intrusion or 'the right to be left alone' and control of information about oneself. The (computer) systems that do the collection and processing of data should therefore be designed with care for privacy. Designing such systems that preserve privacy is a difficult task (if possible at all), in particular when the system is centered around the processing of privacy sensitive data (such as medical information). Fortunately, there is a long history of security principles and legislative work that can be used as a starting point for designing such systems for privacy.

The easiest way to design a privacy preserving system is to not collect, store or process any personal data. However, in practice many computerized systems need to process some personal data. For a large subset of these systems there is no direct explicit need to use personal data, i.e., there are no *functional* requirements to the system to collect, store and process personal data. For example, public transportation systems often use computerized tokens, such as the Oyster system of the London underground or the Dutch OV-chip card, that users have to use to gain access to the public transport system. It can be useful, for example for future planning or optimization purposes, for such systems to collect data about the number of travellers per train. But there is no reason –except for a commercial one– to store the entire travel history for each individual user (as the Dutch system does). Systems that collect personal data for commercial reasons usually do this to be able to provide personalized

(targeted) advertisements or to sell the collected data to other interested parties such as advertisers or insurance companies. Large data processors such as Google and Facebook (but also many less known ones) specialize in this: they are *designed to break privacy* – in particular, users lose control about their own information. They are given an incentive to ‘trade away’ (part of) their privacy (control over personal data) in exchange for small monetary discounts (Groupon) or specific services (Google, Facebook).

Designing for privacy is not limited to only computer systems, some systems such as RFID tags (Juels, 2006), smart phones and the Internet-of-things (Atzori, Iera, & Morabito, 2010) combine physical devices with computer back-ends which leads to all kinds of complications in (privacy preserving) systems design. Other examples such as DNA sequencing have no direct relation with computer systems, but clearly have a privacy impact (and the, privacy sensitive, results of these techniques are often stored in computer systems). For all these systems it is important to design rules and guidelines that enforce the privacy of the users (or subjects) of the system. Section 3 discusses such a system, and its privacy implications, at the border of computer and physical system: the Smart Grid.

This chapter explores the notion of privacy and how to design ‘privacy preserving’ systems: systems that are designed with privacy in mind and systems that can be used to circumvent the large data collectors such as Google and Facebook. Examples of privacy preserving (and breaking) systems, ranging from Smart Meters to Electronic Health records are used to illustrate the main difficulties of designing such systems.

2. Privacy

There is no commonly accepted definition of the concept ‘privacy’. Perhaps this is not surprising since the concept is widely studied in such diverse fields as philosophy, law, social sciences and computer sciences. This section provides a definition of ‘privacy’ that should be acceptable to most. More esoteric -less accepted- notions related to privacy are also discussed.

2.1 Existing relevant definitions, conceptualizations, specifications of privacy

The concept of privacy can be defined in numerous ways and from various perspectives. This chapter discusses the concept of privacy from a philosophical (ontological, ethical) and a legal perspective.

From an ontological perspective it is clear that ‘privacy’ is a social and indeed a cultural (Zakaria et al, 2003; Liu et al 2004) construct: without other people the concept of privacy is meaningless. Privacy is also a right -indeed a fundamental human right (Movios and Krup, 2009), and as such it can be claimed and enforced through legal means. The following three aspects aim to capture the main points associated with the concept of ‘privacy’:

1. freedom from intrusion, the right to be left alone
2. control of information about oneself
3. freedom from surveillance, the right to not be tracked, followed or watched (in one’s own private space)

The first of the above aspects is identical to what Isaiah Berlin called ‘negative liberty’:

“Liberty in the negative sense involves an answer to the question: ‘What is the area within which the subject — a person or group of persons — is or should be left to do or be what he is able to do or be, without interference by other persons’.”(Berlin, 1958)

Negative liberty, and thus also privacy, strives for freedom from external constraints. It deals with relations between people (social!). Individuals typically want to be left alone by larger groups such as organizations and states. In contrast, 'positive liberty' is defined as freedom from 'internal constraints' such as social and culture structures. This is sometimes also explained as the freedom to express oneself as one wants (self-mastery). Privacy can be seen as a necessary precondition for self-expression, and thus for positive liberty, as argued by van den Hoven en Vermaas (2007). In this view privacy is seen as *respect for moral autonomy*, the autonomy to write one's own history and identify with one's own moral choices without "critical gaze, interference of others" (van den Hoven en Vermaas, 2007).

The second and third aspect of privacy, as defined above, are more closely linked to legal notions of privacy. These deal with the control and storing/capturing of information about individuals. Regulations, guidelines and laws such as the EU Data Directive (Birnback, 2008; EU Directive, 1995) and the United States Federal Trade Commission's Fair Information Practice Principles (Annecharico, 2002) try to capture these two aspects in a number of rules, including: i) Transparency (How is data stored/processed?) ii) Purpose (Why is data stored/processed?) iii) Proportionality (Is this necessary for this goal?) iv) Access (What do they know about me, can I change it?) v) Transfer (Who else has access?).

Different countries have different ways of implementing these principles in laws and regulations. For example, the EU has a very strict privacy regulation (the EU Data Protection Directive, 1995), that is enforced 'top-down' for all organizations and citizens in the whole European Union. In contrast, regulations in the United States are typically more sector specific such as HIPAA (HIPAA, 1996) for the health care sector and the Gramm–Leach–Bliley Act (Janger and Schwartz, 2001) for the financial sector. Moreover, the US favors self-regulation, for example the PCI-DSS (PCI, 2009), that is used in the credit card sector. Also note that such laws and regulations are not static (legal) objects, they are continuously being updated, for example a new version of the EU Data Directive (EU Proposal, 2012) has been proposed (also see the next section).

The right to privacy is, at least to a certain degree, relative. One can have a reasonable expectation of privacy in one's own home (see the third aspect above), but not necessarily in public spaces. People that live in the public eye - royalty, celebrities - also have less expectations of privacy in the current, media centered, society. Note that this makes privacy a context-dependent notion.

For privacy, the context of use and control of information is captured in notions such as 'spheres of justice' or 'spheres of access' (van den Hoven, 1999; Nagenborg, 2009) and 'contextual integrity', as used by Ackerman et al (2001) and Nissenbaum (2010). All these notions have in common that they interpret privacy in a local context. The meaning and value of information has a local (possibly cultural) aspect which should be taken into account when analyzing privacy. Nissenbaum in particular understands privacy in terms of context-relative information norms, and distinguishes norms of appropriateness, and norms of distribution. She defines contexts as "structured social settings, characterized by canonical activities, roles, relationships, power structures, norms (or rules), and internal values (goals, ends, purposes)." (Nissenbaum, 2010, p.132- 134). The role of context as it relates to privacy is particularly important when it comes to the use of 'privacy preserving technologies (PETs)', discussed further in Section 3.

The above definition of privacy is, by intension, rather broad. Others have a slightly more narrow definition. For example, the definition given by the Value Sensitive Design (Friedman & Kahn, 2002) approach is: [Privacy] "refers to a claim, an entitlement, or a right of an individual to determine what information about himself or herself can be communicated to others" (Schoeman, 1984). Note that this definition only captures the second aspect of privacy.

One last aspect related to privacy is that of incentives: Large scale socio-technical systems have many stakeholders, each with their own incentives –also with respect to privacy. End users of data processing systems sometimes will be given an incentive to give up some (control over) their privacy in exchange for a monetary discount or service. Many large data processors (Facebook, Google, Groupon) base their business model on this ‘privacy information for something else’ exchange. This issue is also discussed in more detail below.

2.2 Main issues of contention/controversy

While the right to privacy is considered to be a fundamental human right (Movios and Krup, 2009), this right is certainly not absolute. As already mentioned, the right to privacy is less relevant in public spaces or for public figures. It is not clear how far this ‘lack of the right to privacy’ can be stretched: Courts will penalize journalists and others that have gone too far in this respect. These lines are dynamic and are continuously redefined as society changes.

Also, since (the right to) privacy is considered to be a legal construct, governments can implement (and have implemented) various laws and regulations that are in conflict with the right to privacy. For example, phone taps or other surveillance techniques can be legal in certain jurisdictions as long as specific rules are followed or a court has allowed the phone tap. Governments, the proverbial ‘big brothers’, typically don’t respect their own privacy regulations. Depending on the type of government, ranging from open societies to dictatorships, more restrictive and anti-privacy measures are in place. Of course, in practice (at least in open societies) regulations will only allow governments to monitor its citizens as far as is deemed ‘reasonable and necessary’ for law and order purposes. Interpreting what is ‘reasonable and necessary’ monitoring (and other anti-privacy measures) is ultimately decided by the Courts.

In cases where privacy regulations are clearly in place it can still be difficult for citizens to also claim this right. Companies and other organizations are obliged by law (at least in the EU) to inform citizens of all the data they have about them, if so requested. However, in practice most companies don’t reply to such information requests or give very limited and incomplete information at best (Jones and Soltren, 2005; Phelps, Nowak and Ferrell, 2000). So while citizens have the right to control information about them, this right is not actively enforced. A court order can change this, but this is a relatively big hurdle, especially if one considers that hundreds of organizations store (and share!) personal data about citizens.

The newly proposed EU Data Directive (EU Proposal, 2012) tries to remedy this situation by including, amongst others, regulations that enforce disclosure of information about data breaches within 24 hours after the data breach became known, and regulations that enforce the ‘right to be forgotten’. The latter should, for example, enable citizens to force companies (Facebook, Google etc.) to remove all stored data they have about themselves. However, even if this proposal becomes EU law, there are still a number of problems (Rosen, 2010): first of all, the regulation is again difficult to enforce. Companies can claim that they removed all personal data about an individual, but there is no realistic way that this can be verified. Indeed, removing all backup copies (of to-be-removed data) can be a difficult problem in itself. Moreover, there is also the risk that this right can be used to ‘rewrite history’: it’s only a short step from removing information from Facebook to removing information from Wikipedia¹. Note that the context is again important here.

Another related aspect of privacy deals with the perception that people have of, potentially privacy invading, technologies and their use and in how far ‘privacy’ addresses their moral

¹ As an example, consider the case of Wolfgang Werle. Werle has been convicted for murder in Germany. He used German privacy laws to sue Wikipedia to get this information removed from his German Wikipedia page. After winning the case, Werle’s German Wikipedia page no longer exists, but the information is still accessible from, among others, the English and Dutch Wikipedia pages.

worries. Often people are not so much concerned with ‘privacy’ in the sense of being left alone but want to be protected from harm or unfair treatment. Van der Hoven and Vermaas (2007) identified four reasons that often ground calls for privacy: prevention of information-based harm, prevention of informational inequality, prevention of informational injustice, and respect for moral autonomy. In this view people are not primarily concerned about their privacy when they use a system such as Facebook, but rather are concerned about what is done with their personal data, which could harm or discriminate them.

A final point of contention is what actually counts as personal data. In the EU Data Protection Directive, personal data is defined as:

“any information relating to an identified or identifiable natural person (‘data subject’); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity;” (art. 2 a, EU Data Protection Directive, 1995)

This definition is intentionally left very broad, but it is not clear if it also holds for aggregate data. Such data can, in principal, no longer be used to identify a person. Yet it can still be perceived as an invasion of ones privacy when aggregate data is, for example, used to refuse a mobile phone contract based on your address and aggregate data about the credit worthiness of your postal area. Another related issue is that sometimes, aggregate data can be decomposed into personal data (de-aggregation). This is similar to the problem of (de-)anonymizing discussed below.

3. What does it mean to design for Privacy?

Privacy and privacy preserving technologies have been studied for decades in the field of computer science (Feistel, 1973). This section discusses some of the main principles behind these technologies, and how to design new ones.

3.1 Existing approaches and tools

The field of computer security has many adages such as ‘security is not an add-on feature’ that stress that security has to be ‘designed-in’ from the start. The same holds true for privacy. In essence there are three different ways to design a (computer) system that respects the user’s privacy:

1. Never store any personal information
2. Follow very strict (privacy) rules when storing and processing personal data
3. Only store and process anonymized personal data

The first of these rules obviously works and is by far the surest way to design systems that are ‘privacy-proof’. Unfortunately, it is not always desirable or indeed possible to not store or process any personal data. Many organizations and companies need to store some customer data, ranging from banks to tax offices and hospitals.

For systems that need to handle personal information the second rule above applies. There are several rules, guidelines or best practices for designing privacy preserving systems. Most of these are very general and can be traced back to the principles that are formed by the EU Data Directive: transparency; it should be clear what information is stored, purpose; it should be clear for what purpose the personal data is stored, proportionality; only relevant data should be stored, access; the user should know what personal data about them is stored and they

should be able to change errors and transfer; personal data should only be transferred with explicit permission of the user and the user should be able to request a transfer of personal data. Others, such as the PCI-DSS (PCI, 2009), for example, give very detailed guidelines for privacy and security sensitive systems design for a limited domain (in this case that of the credit card industry and its partners such as retailers and banks). Another source of best practices and (security) guidelines for the design of privacy preserving systems is provided by various ISO Standards (Hone and Eloff, 2002). In addition, the ‘Privacy by Design’ approach as advocated by Cavoukian (2009) and others also provides high level guidelines in the form of seven principles for designing privacy preserving systems. Example principles are “Privacy as the Default Setting” and “End-to-End Security – Full Lifecycle Protection” along with the principles (transparency, proportionality) discussed before. The principles of the Privacy by Design approach take as central notion the idea that “data protection needs to be viewed in proactive rather than reactive terms, making privacy by design preventive and not simply remedial” (Cavoukian, 2010). Privacy by design also advocates that data protection should be central in all phases of product life cycles, from initial design to operational use and disposal. The Value Sensitive Design approach to privacy (Friedman, Kahn & Borning, 2006) proposes similar rules, such as informed consent, i.e., give users the option on what information is stored (or not) and transparency, tell users which information is stored about them.

Furthermore, the principles or rules that are formed by the EU Data Directive are themselves technologically neutral. They do not enforce any specific technological solutions. As such they can also be considered as (high level) ‘design principles’. Systems that are designed with these rules and guidelines in mind should thus –in principle– be in compliance with EU privacy laws and (up until a point) respect the privacy of its users. Note that there is a difference between the design and the implementation of a (computer) system. During the implementation phase software bugs are introduced, some of which can be (mis)used to break the system and extract private information. How to implement bug-free computer systems² remains an open research question (Hoare, 2003). This issue is further discussed in the next section.

The third rule (‘only store and process anonymized personal data’) above consists of two different approaches: i) anonymizing tools such as Tor (Dingledine, Mathewson and Syverson, 2004) and Freenet (Clarke et al, 2001) ii) more general, non-technological ways for anonymizing existing data. For example, patient names can be removed from medical data for research, and age information can be reduced to intervals: the age 35 is then represented as falling in the range 30-40. The idea behind this is that a record can no longer be linked to an individual, while the relevant parts of the data can still be used for scientific or other purposes.

Software tools, such as Tor and Freenet, allow users to anonymously browse the web (with Tor) or anonymously share content (Freenet). Such software tools are usually, somewhat misleadingly, called privacy enhancing technologies (PETs). They employ a number of cryptographic techniques and security protocols in order to ensure their goal of anonymous communication. Technically, both systems use the property that numerous users use the system at the same time. In Tor, messages are encrypted and routed along a number of different computers, thereby obscuring the original sender of the message (and thus providing anonymity). Similarly, in Freenet content is stored -in encrypted form- amongst all users of the system. Since users themselves do not have the necessary decryption keys, they do not know what kind of content is stored, by the system, on their own computer. This provides plausible deniability and privacy. The system can at any time retrieve the encrypted content and send it to different Freenet users.

² Or indeed, how to verify the absence of bugs in computer systems.

A relatively new, but promising, technique for designing privacy preserving systems is 'homomorphic encryption' (Gentry, 2009). Homomorphic encryption allows a data processor to process encrypted data, i.e., users could send personal data in encrypted form and get some useful results, for example recommendations of movies that online friends like, back in encrypted form. The original user can then again decrypt the result and use this without revealing any personal data to the data processor. This technique is currently still in its infancy, it does not scale yet to the large amounts of data stored in today's systems. However, if this could be made to work more efficiently the results have the potential to be revolutionary (for privacy preserving systems).

3.2 Comparison and critical evaluation

As mentioned before, by far the easiest way to ensure that a system is privacy preserving is to not store or process any personal data. Of course, in practice, for many systems this will not be possible. Such systems can use the techniques described in the previous section, but these each have their own problems and limitations. The section gives an overview of these issues.

One method for designing privacy preserving systems is to use the various design principles and best practices such as ISO standards, Privacy by Design or the principles behind the EU Data Directive (transparency, purpose, proportionality, access, transfer). However, there are several problems with this. First of all, such rules and principles are typically rather vague and abstract. What does it mean to make a transparent design or to design for proportionality? The principles need to be interpreted and placed in a context when designing a specific –privacy preserving- system. But different people will interpret the principles differently, and while this is useful in a legal setting where lawyers, prosecutors and judges need enough freedom in their own interpretation of a particular situation (context!), this interpretation room is less helpful when one wants to design a system for a *specific* purpose: if several rules/guidelines are interpreted the resulting system might not be privacy preserving because the interpretations might not fit together (are not composable). A more detailed design approach, with less room for interpretation, does not have this problem. Second, if one could agree on a specific, context-dependent, design of a privacy preserving system, then that system still needs to be implemented. Implementation is another phase wherein choices and interpretations are made: system designs can be implemented in infinitely many ways. Moreover, it is very hard –for non-trivial systems- to verify whether an implementation meets its design/specification (Loeckx, Sieber and Stansifer, 1985). This is even more difficult for non-functional requirements such as 'being privacy preserving' or security properties in general (Warnier, 2006).

Another privacy preserving technique is anonymization of data. The idea is that by removing explicit links to individuals the data can be safely processed for, for example, (medical) research purposes. The problem here is that it is very hard to anonymize data in such a way that all links with an individual are removed *and* the resulting anonymized data is still useful for research purposes. Researchers have shown that it is almost always possible to reconstruct links with individuals by using sophisticated statistical methods (Danezis, Diaz and Troncoso, 2007) and by combining multiple databases (Anderson, 2010) that contain personal information. Ultimately, how to address this issue is a trade-off between protecting privacy and advancing research. It suffices to say that even if databases with personal data are anonymized, access to them should remain restricted.

Dedicated software tools that provide anonymity of their users, such as Tor and Freenet, also have some problems. For example, Tor, the tool that allows anonymized communication and browsing over the Internet, is susceptible to an attack whereby, under certain circumstances, the anonymity of the user is no longer guaranteed (Back, Möller and Stiglic, 2001; Evans, Dingledine and Grothoff, 2009). Freenet (and other tools) have similar problems (Douceur,

2002). Note that for such attacks to work, an attacker needs to have access to large resources that in practice are only realistic for intelligence agencies of countries³. However, there are other risks. Configuring such software tools correctly is difficult for the average user, and when the tools are not correctly configured anonymity of the user is no longer guaranteed. And there is always the risk that the computer on which the privacy preserving software runs is infected by a Trojan horse (or other digital pest) that monitors all communication (and knows the identity of the user). This is another example of the importance of context. Such tools can help to protect one's privacy (by providing anonymity), but that protection is never absolute.

In summary, numerous techniques exist for designing privacy preserving systems, each with their own flaws. In practice, the most successful systems are designed for a specific purpose in a specific context. They typically combine several of the techniques described above.

4. Experiences and Examples

Every system that stores or processes personal data has to be designed with privacy in mind. There are too many of such systems to discuss them here in any exhaustive manner. Instead, this section discusses in some detail one large system, the Smart Grid, as an example of what privacy issues arise in complex socio-technical systems and what mechanisms work and do not work in this context. Some examples of other systems that have similar issues are discussed at the end of the section.

In the future power grid, the Smart Grid (Massoud and Wollenberg, 2005), very large numbers of distributed (renewable) energy sources will be connected to the existing grid. These physically distributed generation installations (e.g., gas turbines, micro turbines, fuel cells, solar panels, wind turbines) will be connected to existing infrastructure. Integrated monitoring and control will make it possible to measure the effect on the Grid, for example to measure thermal stress caused by fluctuations in loading, or fast transients due to DC to AC power conversion. Smart metering (McDaniel and McLaughlin, 2009) devices installed with consumers, enable applications such as peak prevention due to demand side management (Gellings and Chamberlin, 1987) and the forming of virtual power stations (Ogston and Brazier, 2009) by groups of consumers that sell their excess power (provided by solar or wind turbines) back into the grid. However, smart meters also store and process privacy sensitive data, and they should be designed with care. Note the importance of context here: in a virtual power station it is crucial that all consumption and production of electricity is carefully registered (using smart meters). However, this information is only stored and processed locally (within the virtual power station) and not shared with utility companies or other parties outside the virtual power station. Thus smart metering itself does not harm one's privacy, only the specific context in which it is used might lead to a privacy violation.

Smart meter data can reveal many things about the members of a household, for example, it is easy to see from the power consumption pattern if somebody is at home, or how many people are a part of a household. More recently, researchers have shown that it is even possible to identify the movie that is being watched in a house, while other electrical appliances are in use, by solely observing the power consumption of the household (Greveler, Justus and Loehr, 2012).

The privacy problems associated with smart metering have led to various outcomes. For example, legislators are -helped by special interest groups- becoming more aware of the

³ For example, the NSA can almost certainly identify users of the TOR network. See: <https://www.eff.org/deeplinks/2012/03/https-and-tor-working-together-protect-your-privacy-and-security-online> (retrieved 3/3/2012)

problem, which has resulted in the blocking of legislation in 2009 by the Dutch Senate that was supposed to handle the mandatory role out of smart meters in the Netherlands (ESMA 2009). The main arguments against the plan were privacy concerns and a lack a choice for citizens if they wanted to participate (Fan, 2011). Electrical power companies have reacted to this by offering several different metering models for citizens, ranging from the old (off-line) system, to smart meters that are under complete control of the power company (Boekema, 2011). Consumers that give more control to the power companies receive a higher discount, in essence trading privacy for money.

That such a trade-off is not necessary is shown by privacy preserving systems that try to serve both the interests of citizens (who, presumably, want privacy) and power companies (who want specific data on electricity use). A number of such privacy preserving systems have been designed. Such systems are based on the techniques discussed in the previous section, such as anonymization (Efthymiou and Kalogridis, 2010) or homomorphic encryption (Garcia and Jacobs, 2011; Kursawe et al, 2011). Unfortunately, most of these systems are currently not operational. This is partly because of implementation issues, but also because of incentives of power companies and end users. Power companies can make (more) money by offering new services based on user's power consumption data or by selling (aggregated) data to governments and other organizations and end users still do not ask for privacy and are willing to trade privacy for small monetary discounts. This shows again that, in essence, the specific context determines the success of privacy preserving technologies: if someone can make money of privacy sensitive data, it will usually happen (also see Facebook and Google). Legislation can help in such cases, but lack of enforcement remains a major issue.

Other examples of complex socio-technical systems that have similar privacy issues are electronic patient records in the health sector (Barrows and Clayton, 1996; van 't Noordende, 2010), public transport systems (Winters, 2004; Garcia et al, 2008), electronic criminal records (Brazier et al 2004; Warnier et al 2008) and electronic social networks (Gross and Acquisti, 2005; Rosenblum, 2007). All these systems have in common that (i) they store their information in digital form, (ii) they operate on the scale of countries or bigger and (iii) different stakeholders have different incentives, roles and interest in the system, in particular with regards to privacy. The first two points ensure that the systems can process more and more data automatically at ever growing scales, which leads to ever more complex systems with more stakeholders (more organizations, countries and people can become involved). This growing complexity is difficult enough to manage, but if the growing number of stakeholders, with different incentives (the context), is not taken into account more and more of these systems will ultimately (inevitably!) fail to protect the privacy of its users.

5. Open issues and future work

One major (unsolved) issue in the design of privacy preserving systems is that such systems are 'dual-use' (Atlas, and Dando, 2006): They can be used to protect the privacy of citizens and dissidents, but they can also be used for illegal purposes such as terrorism and the distribution of child pornography. As the Freenet faq⁴ states:

“What about child porn, offensive content or terrorism?”

While most people wish that child pornography and terrorism did not exist, humanity should not be deprived of their freedom to communicate just because of how a very small number of people might use that freedom.”

⁴ <http://freenetproject.org/faq.html#childporn> (retrieved 3/3/2012)

This is a serious problem that has no realistic solution, but is too important to ignore (as the Freenet system does). Some privacy preserving systems use key-escrow schemes (Denning and Branstad, 1996) for this: basically, the system allows the use of a master key that can 'open' all encryption used in the system (and thus revealing the identity of criminal users). But it is unclear who should have access to the master key: the government? The United Nations? And if (when) it becomes known that such a key escrow scheme exists, nobody want to use the system anymore, as, for example, the Clipper chip has shown (Fromkin, 1995).

There are good guidelines and methodologies for the design of privacy preserving systems, but there is still a lot of work to be done for the verification and validation of such systems: how do we know that a particular system indeed has the (privacy) properties we want? This remains an open research question.

6. Conclusions

The multi-faceted aspect of the concept privacy, with multiple stakeholders (with their own incentives) makes it difficult to design privacy preserving systems. In general, 'there is no golden bullet', a 'one-size-fits-all' solution to designing privacy preserving systems. The particular context of the system needs to be taken into account. Even when new techniques, such as homomorphic encryption, become available, other (non-technical) issues such as context and incentives will at least be as important (if not more so).

7. Cross References

Design for the value of trust, Design for the Values of accountability and transparency, Design for values in Information and Communication Technology

8. References

- Ackerman, M., Darrell, T and Weitzner, D (2001) Privacy in context. *Human Computer Interaction* 16:167-176
- Anderson, R.J. (2010) *Security Engineering: A guide to building dependable distributed systems*, Wiley.
- Annecharico, D. (2002) Notes & Comments: V. Privacy after GLBA: Online Transactions: Squaring the Gramm-Leach-Bliley Act Privacy Provisions With the FTC Fair Information Practice Principles. *NC Banking Inst.* 6, 637-695.
- Atlas, R.M. and Dando, M. (2006) The dual-use dilemma for the life sciences: perspectives, conundrums, and global solutions, *Biosecurity and bioterrorism: biodefense strategy, practice, and science*, 4(3):276-286.
- Atzori, L., Iera, A., & Morabito, G. (2010) The internet of things: A survey. *Computer Networks*, 54(15):2787-2805.
- Back, A., Möller, U. and Stiglic, A. (2001) Traffic analysis attacks and trade-offs in anonymity providing systems. In: Katzenbeisser, S., & Sadeghi, A. R. (ed) *Information Hiding*, Springer, p 245-257
- Barrows Jr, R.C. and Clayton, P.D. (1996) Privacy, confidentiality, and electronic medical records, *Journal of the American Medical Informatics Association* 3(2):139-148
- Berlin, I. (1958). *Two concepts of liberty*, Clarendon Press Oxford
- Birnhack, M. (2008). The EU Data Protection Directive: An engine of a global regime. *Computer Law & Security Report* 24(6):508-520.
- Boekema, J. (2011) Assessment of the implementation regulations for Smart Meters, TNO Technical Report, TNO-RPT-DTS-2011-00463-E

- Brazier, F. M. T., Oskamp, A., Prins, J. E. J., Schellekens, M. H. M. and Wijngaards, N. J. E., (2004) Anonymity and Software Agents: An Interdisciplinary Challenge. *AI & Law*, 1-2(12):137-157
- Cavoukian, A. (2009) *Privacy by Design*. Ottawa: IPC
- Cavoukian, A. (2010) Privacy by Design: The Definitive workshop. *Identity in the Information Society* 3(2):121-126
- Clarke, I., Sandberg, O., Wiley, B. and Hong, T (2001). Freenet: A distributed anonymous information storage and retrieval system. In: Federrath, H. (ed) *Designing Privacy Enhancing Technologies*, Springer, p 46-66
- Danezis, G., Diaz, C. and Troncoso, C. (2007) Two-sided statistical disclosure attack, In: *Proceedings of the 7th international conference on Privacy enhancing technologies*, Springer, p 30-44
- Denning, D.E. and Branstad, D.K. (1996) Key Escrow Encryption Systems, *Communications of the ACM*, 39(3):35.
- Dingledine, R., Mathewson, N. and Syverson, P. (2004) Tor: The second-generation onion router, *Proceedings of the 13th conference on USENIX Security Symposium-Volume 13*, page 21.
- Douceur, J. (2002) The Sybil attack, In: *Peer-to-peer Systems*, Springer, p 251-260,
- Efthymiou, C. and Kalogridis, G (2010) Smart grid privacy via anonymization of smart metering data, in *First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, p 238-243.
- EU Data Protection Directive (1995). Directive 95/46/EC of the European Parliament and of 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.
- EU Proposal (2012). Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation).
- ESMA (2009) *Annual Report on the Progress in Smart Metering, Version 2.0*
- Evans, N.S., Dingledine, R. and Grothoff, C. (2009) A practical congestion attack on Tor using long paths, In: *Proceedings of the 18th conference on USENIX security symposium*, p 33-50.
- Fan, Z. and Kulkarni, P. and Gormus, S. and Efthymiou, C. and Kalogridis, G. and Sooriyabandara, M. and Zhu, Z. and Lambotharan, S. and Chin, W. (2011) Smart Grid Communications: Overview of Research Challenges, Solutions, and Standardization Activities, *Communications Surveys & Tutorials*, IEEE 99:1-18
- Feistel, H. (1973). Cryptography and computer privacy, *Scientific American*, 228(5):15-23.
- Friedman, B., & Kahn Jr, P. H. (2002, January). Human values, ethics, and design. In: *The human-computer interaction handbook* L. Erlbaum Associates Inc. p 1177-1201
- Friedman, B., Kahn Jr, P. H., & Borning, A. (2006). Value sensitive design and information systems. *Human-computer interaction in management information systems: Foundations*, 5:348-372.
- Froomkin, A.M. (1995) The metaphor is the key: cryptography, the clipper chip, and the constitution, *University of Pennsylvania Law Review* 143(3):709-897.
- Garcia, F.D. and Jacobs, B.P.F. (2011) Privacy-friendly Energy-metering via Homomorphic Encryption. In *6th Workshop on Security and Trust Management (STM 2010) Lecture Notes in Computer Science*, Volume 6710, Springer, p 226-238,
- Garcia, F.D., de Koning Gans, G., Muijers, R. van Rossum, P., Verdult, R, Wichers Schreur, R. and Jacobs, B.P.F. (2008) Dismantling MIFARE Classic. In S. Jajodia, and J. Lopez (ed) *13th European Symposium on Research in Computer Security (ESORICS 2008)*. *Lecture Notes in Computer Science*, Vol. 5283, Springer, p 97-114
- Gellings, C.W. and Chamberlin, J.H. (1987). *Demand-side management: concepts and methods*, The Fairmont Press Inc., Lilburn, GA
- Gentry, C (2009). Fully homomorphic encryption using ideal lattices, In: *Proceedings of the 41st annual ACM symposium on Theory of computing*, ACM, p169-178,
- Greveler, U., Justus, B. and Loehr, D. (2012) Multimedia Content Identification Through Smart Meter Power Usage Proles, In: *Computers, Privacy and Data Protection*, Springer
- Gross, R. and Acquisti, A. (2005) Information revelation and privacy in online social networks, In: *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, ACM, p 71-80,
- HIPAA (1996). *Health Insurance Portability and Accountability Act of 1996*.
- Hoare, T. (2003) The verifying compiler: A grand challenge for computing research. In: *Proceedings of the 12th international conference on Compiler construction*, Springer, p 262-272

- Hone, K. and Eloff, J.H.P. (2002) Information security policy--what do international information security standards say? *Computers & Security* 21(5):402-409
- Hoven, M.J. van den (1999) Privacy or informational injustice? In: Pourcia L (ed) *Ethics and information in the twenty-first century*, Purdue University Press, p 140–150
- Van den Hoven, J., & Vermaas, P. E. (2007). Nano-technology and privacy: on continuous surveillance outside the panopticon. *Journal of Medicine and Philosophy*, 32(3):283-297.
- Janger, E.J. and Schwartz, P.M. (2001). The Gramm-Leach-Bliley Act, Information Privacy, and the Limits of Default Rules, *Minnesota Law Review* 86.
- Jones, H. and Soltren, H (2005). Facebook: Threats to privacy, *Social Science Research*, p 1-76.
- Juels, A. (2006). RFID security and privacy: A research survey. *Selected Areas in Communications, IEEE Journal on*, 24(2):381-394.
- Kursawe, K., Danezis, G., & Kohlweiss, M. (2011). Privacy-friendly aggregation for the smart-grid. In: *Privacy Enhancing Technologies*, Springer, p 175-191.
- Liu, C. and Marchewka, J.T. and Ku, C. (2004) American and Taiwanese perceptions concerning privacy, trust, and behavioral intentions in electronic commerce, *Journal of Global Information Management*, p 18-40.
- Loeckx, J. and Sieber, K. and Stansifer, R.D. (1985) *The foundations of program verification*, John Wiley & Sons
- Massoud, S.A. and Wollenberg, B (2005) Toward a smart grid: power delivery for the 21st century. *Power and Energy Magazine, IEEE*, 3(5):34–41.
- McDaniel, P. and McLaughlin, S. (2009) Security and privacy challenges in the smart grid. *IEEE Security & Privacy* 7(3):75-77.
- Movius, L. and N. Krup (2009). U.S. and EU Privacy Policy: Comparison of Regulatory Approaches. *International Journal of Communication*, p 169–187.
- Nagenborg, M. (2009) Designing spheres of informational justice. *Ethics and Information Technology* 11:175–179
- Nissenbaum, H. (2010) *Privacy in Context*. Stanford University Press
- Orgston, E.F.Y. and Brazier, F.M.T. (2009) Apportionment of control in virtual power stations, *Proceedings of the international conference on infrastructure systems and services 2009: developing 21st century infrastructure networks, IEEE computer society*, p 1-6
- PCI (2009). PCI Security Standards Council, Payment Card Industry (PCI) Data Security Standard - Requirements and Security Assessment Procedures version 1.2.
- Phelps, J. and Nowak, G. and Ferrell, E (2000). Privacy concerns and consumer willingness to provide personal information, *Journal of Public Policy & Marketing*.
- Rosen, J. (2010) The End of Forgetting, *The New York Times Magazine*, July 25.
- Rosenblum, D. (2007) What anyone can know: The privacy risks of social networking sites, *IEEE Security & Privacy, IEEE Computer Society*. p 40-49,
- Schoeman, F. D. (Ed.). (1984). *Philosophical dimensions of privacy: An anthology*. Cambridge University Press.
- van 't Noordende, G. (2010) Security in the Dutch Electronic Patient Record System, *2nd ACM Workshop on Security and Privacy in Medical and Home-Care Systems (SPIMACS)*, p 21-31.
- Warnier, M.E. (2006) *Language Based Security for Java and JML*, PhD Thesis, Radboud University Nijmegen.
- Warnier, M.E., Brazier, F. M. T. and Oskamp, A., (2008) Security of Distributed Digital Criminal Dossiers, in: *Journal of Software*, 3(3):21-29, Academy Publisher.
- Warren, S.D. and Brandeis, L.D. (1890) The Right to Privacy, *Harvard Law review* 4(5):193-220.
- Winters, N. (2004) Personal privacy and popular ubiquitous technology, In: *Proceedings of Ubiconf*
- Zakaria, N. and Stanton, J.M. and Sarkar-Barney, S.T.M. (2003) Designing and implementing culturally-sensitive IT applications: The interaction of culture values and privacy issues in the Middle East, *Information Technology & People* 16(1):49-75.