

Privacy Design Strategies

Jaap-Henk Hoepman

► **To cite this version:**

Jaap-Henk Hoepman. Privacy Design Strategies. Nora Cuppens-Boulahia; Frédéric Cuppens; Sushil Jajodia; Anas Abou El Kalam; Thierry Sans. 29th IFIP International Information Security Conference (SEC), Jun 2014, Marrakech, Morocco. Springer, IFIP Advances in Information and Communication Technology, AICT-428, pp.446-459, 2014, ICT Systems Security and Privacy Protection. <10.1007/978-3-642-55415-5_38>. <hal-01370395>

HAL Id: hal-01370395

<https://hal.inria.fr/hal-01370395>

Submitted on 22 Sep 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Privacy Design Strategies (extended abstract)*

Jaap-Henk Hoepman¹

Radboud University Nijmegen
P.O. Box 9010, 6500 GL Nijmegen, the Netherlands
jhh@cs.ru.nl.

Abstract. In this paper we define the notion of a privacy design strategy. These strategies help IT architects to support privacy by design early in the software development life cycle, during concept development and analysis. Using current data protection legislation as point of departure we derive the following eight privacy design strategies: MINIMISE, HIDE, SEPARATE, AGGREGATE, INFORM, CONTROL, ENFORCE, and DEMONSTRATE. The strategies also provide a useful classification of privacy design patterns and the underlying privacy enhancing technologies. We therefore believe that these privacy design strategies are not only useful when designing privacy friendly systems, but also helpful when evaluating the privacy impact of existing IT systems.

1 Introduction

Privacy by design [5] is a system design philosophy that aims to improve the overall privacy¹ friendliness of IT systems. Point of departure is the observation that privacy (like security) is a core property of a system that is heavily influenced by the underlying system design. As a consequence, privacy protection cannot be implemented as an add-on. Privacy must be addressed from the outset instead. The fundamental principle of privacy by design is, therefore, that privacy requirements must be addressed throughout the full system development process. In other words starting when the initial concepts and ideas for a new system are drafted, up to and including the final implementation of that system. Privacy by design is gaining importance. For example, the proposal for a new European data protection regulation [10] explicitly requires data protection by design and by default. It is therefore crucial to support developers in satisfying these requirements with practical tools and guidelines.

As explained in Section 2, an important design methodology is the application of so called software design patterns. These design patterns refine the system architecture to achieve certain functional requirements within a given set of constraints.

* This research is supported by the research program Sentinels (www.sentinels.nl) as project 'Revocable Privacy' (10532). Sentinels is being financed by Technology Foundation STW, the Netherlands Organization for Scientific Research (NWO), and the Dutch Ministry of Economic Affairs. This research was (partially) conducted within the Privacy and Identity Lab (PI.lab, www.pilab.nl).

Version: Fri Mar 14 20:32:18 2014 +0100 / ifip-sec-2014-final

¹ In this paper we focus on data protection, and treat privacy and data-protection as synonyms.

During software development the availability of practical methods to protect privacy is high during actual implementation, but low when starting the project. Numerous privacy enhancing technologies (PETs) exists that can be applied more or less 'off the shelf'. Before that implementation stage, privacy design patterns can be used during system design. Significantly less design patterns exist compared to PETs, however. And at the start of the project, during the concept development and analysis phases, the developer stands basically empty handed.

This paper aims to close this gap [13,26]. Design patterns do not necessarily play a role in the earlier, concept development and analysis, phases of the software development cycle. The main reason is that such design patterns are already quite detailed in nature, and more geared towards solving an implementation problem. To guide the development team in the earlier stages, we define the notion of a privacy design strategy. Because these strategies describe fundamental, more strategic, approaches to protecting privacy, they enable the IT developer to make well founded choices during the concept development and analysis phase as well. These choices have a huge impact on the overall privacy protection properties of the final system.

The privacy design strategies developed in this paper are derived from existing privacy principles and data protection laws. These are described in section 3. We focus on the principles and laws on which the design of an IT system has a potential impact. By taking an abstract information storage model of an IT system as a point of departure, these legal principles are translated to a context more relevant for the IT developer in section 4. This leads us to define the following privacy design strategies: MINIMISE, HIDE, SEPARATE, AGGREGATE, INFORM, CONTROL, ENFORCE and DEMONSTRATE. They are described in detail in section 5.

We believe these strategies help to support privacy by design throughout the full software development life cycle, even before the design phase. It makes explicit which high level decisions can be made to protect privacy, when the first concepts for a new information system are drafted. The strategies also provide a useful classification of privacy design patterns and the underlying privacy enhancing technologies. We therefore believe that these privacy design strategies are not only useful when designing privacy friendly systems, but that they also provide a starting point for evaluating the privacy impact of existing information systems.

2 Software development

Software architecture encompasses the set of significant decisions about the organisation of a software system², including the selection of the structural elements and their interfaces by which a system is composed, the behaviour as specified in collaborations among those elements, the composition of these structural and behavioural elements into larger subsystem, and the architectural style that guides this organisation.

Typically, the development of a software system proceeds in six phases: concept development, analysis, design, implementation, testing and evaluation. In fact,

² Based on an original definition by Mary Shaw, expanded in 1995 by Grady Booch, Kurt Bittner, Philippe Kruchten and Rich Reitman as reported in [19].

these phases are often considered a cycle, where after evaluation a new iteration starts by updating the concept as appropriate. In this paper we distinguish design strategies (defined in this paper), design patterns and concrete (privacy enhancing) technologies as tools to support the decisions to be made in each of these phases. The design strategies support the concept development and analysis phases, the design patterns are applicable during the design phase, and the privacy enhancing technologies are useful during the implementation phase.

2.1 Design patterns

The concept of a *design pattern* is a useful vehicle for making design decisions about the organisation of a software system. A design pattern

“provides a scheme for refining the subsystems or components of a software system, or the relationships between them. It describes a commonly recurring structure of communicating components that solves a general design problem within a particular context.” [3]

Typically, the description [11] of a design pattern contains at least its name, purpose, context (the situations in which it applies), implementation (its structure, components and their relationships), and the consequences (its results, side effects and trade-offs when applied). Many design patterns exist, at varying levels of abstraction.

A classical software design pattern is the *Model-View-Controller*³, that separates the representation of the data (the model) from the way it is represented towards the user (the view) and how the user can modify that data (using the controller). Few privacy design patterns have been *explicitly* described as such to date. We are aware of the work of Hafiz [14,15], Pearson [23,22], van Rest *et al.* [30], and a recent initiative of the UC Berkeley School of Information⁴. Many more *implicit* privacy design patterns exist though, although they have never been described as such. We will encounter them in our discussion of the patterns corresponding to the privacy design strategies we develop below.

2.2 Design strategies

Because certain design patterns have a higher level of abstraction than others, some authors also distinguish *architecture patterns*, that

“express a fundamental structural organisation or schema for software systems. They provide a set of predefined subsystems, specify their responsibilities, and include rules and guidelines for organising the relationships between them.”⁵

³ Originally formulated in the late 1970s by Trygve Reenskaug at Xerox PARC, as part of the Smalltalk system.

⁴ <http://privacypatterns.org/>

⁵ See <http://best-practice-software-engineering.ifs.tuwien.ac.at/patterns.html>, and The Open Group Architecture Framework (TOGAF) <http://pubs.opengroup.org/architecture/togaf8-doc/arch/chap28.html>

The *Model-View-Controller* pattern cited above is sometimes considered such an architecture pattern. The distinction between an architecture pattern and a design pattern is not always easily made, however. Moreover, there are even more general principles that guide the system architecture. We choose, therefore, to express such higher level abstractions in terms of *design strategies*. We define this as follows.

A design strategy describes a fundamental approach to achieve a certain design goal. It favours certain structural organisations or schemes over others. It has certain properties that allow it to be distinguished from other (fundamental) approaches that achieve the same goal.

Whether something classifies as a strategy very much depends on the universe of discourse, and in particular on the exact goal the strategy aims to achieve. A *privacy design strategy* is a design strategy that achieves (some level of) privacy protection as its goal.

Design strategies do not necessarily impose a specific structure on the system although they certainly limit the possible structural realisations of it. Therefore, they are also applicable during the concept development and analysis phase of the development cycle⁶.

2.3 Privacy enhancing technologies

Privacy Enhancing Technologies (PETs) are better known, and much more studied. Borking and Blarckom *et al.* [1,29] define them as follows.

“Privacy-Enhancing Technologies is a system of ICT measures protecting informational privacy by eliminating or minimising personal data thereby preventing unnecessary or unwanted processing of personal data, without the loss of the functionality of the information system.”

This definition was later adopted almost literally by the European Commission [8].

In principle, PETs are used to implement a certain privacy design pattern with concrete technology. For example, both ‘Idemix’ [4] and ‘u-prove’ [2] are privacy enhancing technologies implementing the (implicit) design pattern *anonymous credentials*. There are many more examples of privacy enhancing technologies, like ‘cut-and-choose’ techniques [7], ‘onion routing’ [6] to name but a few.

3 The foundations of data protection

We aim to derive privacy design strategies from existing data protection laws and privacy frameworks. We therefore briefly summarise those here.

In the European Union, the legal right to privacy is based on Article 8 of the European Convention of Human Rights of 1950. In the context of data protection,

⁶ We note that the notion of a privacy design strategy should not be confused with the foundational principles of Cavoukian [5] or the concept of a privacy principle from the ISO 29100 Privacy framework [17].

this right has been made explicit in the 1995 data protection directive [9], which is based on the privacy guidelines of the Organisation of Economic Co-Operation and Development (OECD) from 1980 [21].

3.1 The OECD guidelines

The OECD guidelines, of which the US fair information practices (FIPs) [28] — *notice, choice, access* and *security* — are a subset, define the following principles.

- The collection of personal data is lawful, limited, and happens with the knowledge or consent of the data subject (*Collection Limitation*).
- Personal data should be relevant to the purposes for which they are to be used, and be accurate, complete and kept up-to-date (*Data Quality*).
- The purposes of the collection must be specified upfront (*Purpose Specification*), and the use of the data after collection is limited to that purpose (*Use Limitation*).
- Personal data must be adequately protected (*Security Safeguards*).
- The nature and extent of the data processing and the controller responsible must be readily available (*Openness*).
- Individuals have the right to view, erase, rectify, complete or amend personal data stored that relates to him (*Individual Participation*).
- A data controller must be accountable for complying with these principles (*Accountability*).

3.2 Data protection in Europe

The OECD principles correspond roughly to the main provisions in the European data protection Directive of 1995 [9]. For example, Article 6 states that personal data must be processed fairly and lawfully, must be collected for a specified purpose, and must not be further processed in a way incompatible with those purposes. Moreover the data must be adequate, relevant, and not excessive. It must be accurate and up to date, and kept no longer than necessary. These provisions express a need for purpose limitation, data minimisation, and data quality.

Other articles of the Directive deal with transparency and user choice. For example, article 7 requires unambiguous consent from the data subject, while article 10 and 11 require data controllers to inform data subjects about the processing of personal data. Article 12 gives data subjects the right to review and correct the personal data that is being processed about them. Finally, security as a means to protect privacy is addressed by article 17, that mandates adequate security of processing.

We note that the European data protection directive covers many more aspects, that are however less relevant for the discussion in this paper. The directive is currently under review. A proposal for a regulation to replace it has been published [10], and an amendment was recently adopted by the European Parliament. This regulation is still in flux and under heavy debate, but it contains the following rights and obligations that are relevant for our discussion in this paper.

- A controller must implement data protection by design and by default (art. 23).

- A controller must be able to demonstrate compliance with the regulation (art. 5, and also art. 22).
- Data subjects have the right to be forgotten and to erasure (art. 17).
- Data subjects have the right to data portability, allowing them ‘to obtain from the controller a copy of data undergoing processing in an electronic and structured format which is commonly used and allows for further use by the data subject’ (art. 18).
- A data controller has the duty to issue a notification whenever a personal data breach occurs (art. 31 and 32).

3.3 The ISO 29100 Privacy Framework perspective

In the full paper [16] we also include the ISO 29100 Privacy framework [17] in our analysis, but due to space constraints this is omitted from this extended abstract.

3.4 Summary of requirements

Not every legal requirement can be met by designing an IT system in a specific way. Legitimacy of processing is a good example. If the processing is illegitimate, then it will be illegitimate irrespective of the design of the system. We therefore focus our effort on studying aspects on which the design of an IT system has a potential impact. These are summarised in the list below.

- Purpose limitation (comprising both specification of the purpose and limiting the use to that stated purpose).
- Data minimisation.
- Data quality.
- Transparency (Openness in OECD terms).
- Data subject rights (in terms of consent, and the right to view, erase, and rectify personal data).
- The right to be forgotten.
- Adequate protection (Security Safeguards in OECD terms).
- Data portability.
- Data breach notifications.
- Accountability and (provable) compliance.

These principles must be covered by the privacy design strategies that we will derive next. Whether this is indeed the case is analysed in section 6.

4 Deriving privacy design strategies

A natural starting point to derive privacy preserving strategies is to look at when and how privacy is violated, and then consider how these violations can be prevented. Solove’s taxonomy [25], for example, identifies four basic groups of activities that affect privacy: information collection, information processing, information dissemination and invasions. This is in fact similar to the distinction made between

data transfer, storage and processing by Spiekermann and Cranor [26]. This general subdivision inspired us to look at IT systems at a higher level of abstraction to determine where and how privacy violations can be prevented.

In doing so, we can view an IT system as an information storage system (i.e., database system) system. Many of today's systems, like classical business or government administration systems, are database systems. The same holds for social networks. Current data protection legislation [9] is pretty much written with such a database model in mind. In a database, information about individuals is stored in one or more tables. Each table stores a fixed set of attributes for an individual. The columns in the table represent this fixed set of attributes. A row is added for each new individual about whom a record needs to be stored. Sometimes, data is not stored at the level of individual persons, but is instead aggregated based on certain relevant group properties (like postal code).

Within the legal framework described in section 3, the collection of personal data should be proportional to the purpose for which it is collected, and this purpose should not be achievable through other, less invasive means. In practice, this means that data collection should be *minimised*. This can be achieved by not storing individual rows in a database table for each and every individual. Also the collection of attributes stored should correspond to the purpose, leading to fewer columns being stored. Data collected for one purpose should be stored *separately* from data stored for another purpose. Linking of these database tables should not be easy. When data about specific individuals is not necessary for the purpose, only *aggregate* data should be stored. Personal data should be properly protected, and *hidden* from other parties. A data subject should be *informed* about the fact that data about her is being processed, and she should be able to request modifications and corrections where appropriate. In fact the underlying principle of information self-determination [31] dictates the she should be in *control*. Finally, the collection and processing of personal data should be done in accordance to a privacy policy, that should be actively *enforced*. The current proposal for the revision of the European privacy directive (into a regulation) also stresses the fact that data controllers should be able to *demonstrate compliance* with data protection legislation. The data controller has the burden of proof with respect to compliance and must, for example, run and document a privacy impact assessment (PIA).

Given this analysis from the legal point of view, we distinguish the following eight privacy design strategies: MINIMISE, SEPARATE, AGGREGATE, HIDE, INFORM, CONTROL, ENFORCE and DEMONSTRATE. A graphical representation of these strategies, when applied to a database system, is given in Figure 1.

5 The eight privacy design strategies

We will now proceed to describe these eight strategies in detail. We have grouped the strategies into two classes: data-oriented strategies and process-oriented strategies. The first class roughly corresponds to the privacy-by-architecture approach identified by Spiekermann and Cranor [26], whereas the process-oriented strategies more-or-less cover their privacy-by-policy approach.




Fig. 1. The database metaphor of the eight privacy design strategies.

5.1 Data oriented strategies

Strategy #1: MINIMISE The most basic privacy design strategy is MINIMISE, which states that

The amount of personal data that is processed⁷ should be restricted to the minimal amount possible.

This strategy is extensively discussed by Gürses *et al.* [13]. By ensuring that no, or no unnecessary, data is collected, the possible privacy impact of a system is limited. Applying the MINIMISE strategy means one has to answer whether the processing of personal data is proportional (with respect to the purpose) and whether no other, less invasive, means exist to achieve the same purpose. The decision to collect personal data can be made at design time and at run time, and can take various forms. For example, one can decide not to collect any information about a particular data subject at all. Alternatively, one can decide to collect only a limited set of attributes.

Design patterns Common design patterns that implements this strategy are *select before you collect* [18], *anonymisation* and *use pseudonyms* [24].

Strategy #2: HIDE The second design strategy, HIDE, states that

Any personal data, and their interrelationships, should be hidden from plain view.

⁷ For brevity, and in line with Article 2 of the European directive [9], we use data processing to include the collection, storage and dissemination of that data as well.

The rationale behind this strategy is that by hiding personal data from plain view, it cannot easily be abused. The strategy does not directly say from whom the data should be hidden. And this depends on the specific context in which this strategy is applied. In certain cases, where the strategy is used to hide information that spontaneously emerges from the use of a system (communication patterns for example), the intent is to hide the information from anybody. In other cases, where information is collected, stored or processed legitimately by one party, the intent is to hide the information from any other party. In this case, the strategy corresponds to ensuring confidentiality.

The HIDE strategy is important, and often overlooked. In the past, many systems have been designed using innocuous identifiers that later turned out to be privacy nightmares. Examples are identifiers on RFID tags, wireless network identifiers, and even IP addresses. The HIDE strategy forces one to rethink the use of such identifiers. In essence, the HIDE strategy aims to achieve unlinkability and unobservability [24]. Unlinkability in this context ensures that two events cannot be related to one another (where events can be understood to include data subjects doing something, as well as data items that occur as the result of an event).

Design patterns The design patterns that belong to the HIDE strategy are a mixed bag. One of them is the use of *encryption* of data (when stored, or when in transit). Other examples are *mix networks* [6] to hide traffic patterns [6], or techniques to unlink certain related events like *attribute based credentials* [4], *anonymisation* and the use of *pseudonyms*. Note that the latter two patterns also belong to the MINIMISE strategy.

Strategy #3: SEPARATE The third design strategy, SEPARATE, states that

Personal data should be processed in a distributed fashion, in separate compartments whenever possible.

By separating the processing or storage of several sources of personal data that belong to the same person, complete profiles of one person cannot be made. Moreover, separation is a good method to achieve purpose limitation. The strategy of separation calls for distributed processing instead of centralised solutions. In particular, data from separate sources should be stored in separate databases, and these databases should not be linked. Data should be processed locally whenever possible, and stored locally if feasible as well. Database tables should be split when possible. Rows in these tables should be hard to link to each other, for example by removing any identifiers, or using table specific pseudonyms.

These days, with an emphasis on centralised web based services this strategy is often disregarded. However, the privacy guarantees offered by peer-to-peer networks are considerable. Decentralised social networks like Diaspora⁸ are inherently more privacy friendly than centralised approaches like Facebook and Google+.

Design patterns No specific design patterns for this strategy are known.

⁸ <http://diasporafoundation.org/>

Strategy #4: AGGREGATE The fourth design pattern, AGGREGATE, states that

Personal data should be processed at the highest level of aggregation and with the least possible detail in which it is (still) useful.

Aggregation of information over groups of attributes or groups of individuals, restricts the amount of detail in the personal data that remains. This data therefore becomes less sensitive. When the information is sufficiently coarse grained, and the size of the group over which it is aggregated is sufficiently large, little information can be attributed to a single person, thus protecting its privacy.

Design patterns Examples of design patterns that belong to this strategy are the following: *aggregation over time* (used in smart metering), *dynamic location granularity* (used in location based services), and *k-anonymity* [27].

5.2 Process oriented strategies

Strategy #5: INFORM The INFORM strategy corresponds to the important notion of transparency:

Data subjects should be adequately informed whenever personal data is processed.

Whenever data subjects use a system, they should be informed about which information is processed, for what purpose, and by which means. This includes information about the ways the information is protected, and being transparent about the security of the system. Providing access to clear design documentation is also a good practice. Data subjects should also be informed about third parties with which information is shared. And data subjects should be informed about their data access rights and how to exercise them.

Design patterns A possible design patterns in this category is the *Platform for Privacy Preferences (P3P)*⁹. *Data breach notifications* are also a design pattern in this category. Finally, Graf *et al.* [12] provide an interesting collection of privacy design patterns for informing the user from the Human Computer Interfacing perspective.

Strategy #6: CONTROL The control strategy states that

Data subjects should be provided agency over the processing of their personal data.

The CONTROL strategy is in fact an important counterpart to the INFORM strategy. Without reasonable means of controlling the use of one's personal data, there is little use in informing a data subject about the fact that personal data is collected. Of course the converse also holds: without proper information, there is little use in asking consent. Data protection legislation often gives the data subject the right to

⁹ <http://www.w3.org/P3P/>

view, update and even ask the deletion of personal data collected about her. This strategy underlines this fact, and design patterns in this class give users the tools to exert their data protection rights.

CONTROL goes beyond the strict implementation of data protection rights, however. It also governs the means by which users can decide whether to use a certain system, and the way they control what kind of information is processed about them. In the context of social networks, for example, the ease with which the user can update his privacy settings through the user interface determines the level of control to a large extent. So user interaction design is an important factor as well. Moreover, by providing users direct control over their own personal data, they are more likely to correct errors. As a result the quality of personal data that is processed may increase.

Design patterns We are not aware of specific design patterns that fit this strategy.

Strategy #7: ENFORCE The seventh strategy, ENFORCE, states:

A privacy policy compatible with legal requirements should be in place and should be enforced.

The ENFORCE strategy ensures that a privacy policy is in place. This is an important step in ensuring that a system respects privacy during its operation. Of course the actual level of privacy protection depends on the actual policy. At the very least it should be compatible with legal requirements. As a result, purpose limitation is covered by this strategy as well. More importantly though, the policy should be enforced. This implies, at the very least, that proper technical protection mechanisms are in place that prevent violations of the privacy policy. Moreover, appropriate governance structures to enforce that policy must also be established.

Design patterns Access control is an example of a design patterns that implement this strategy. Another example are *sticky policies* and privacy rights management: a form of digital rights management involving licenses to personal data.

Strategy #8: DEMONSTRATE The final strategy, DEMONSTRATE, requires a data controller to

Be able to demonstrate compliance with the privacy policy and any applicable legal requirements.

This strategy goes one step further than the ENFORCE strategy in that it requires the data controller to prove that it is in control. This is explicitly required in the new draft EU privacy regulation [10]. In particular this requires the data controller to be able to show how the privacy policy is effectively implemented within the IT system. In case of complaints or problems, she should immediately be able to determine the extent of any possible privacy breaches.

	Purpose limitation	Data minimisation	Data quality	Transparency	Data subject rights	The right to be forgotten	Adequate protection	Data portability	Data breach notification	(Provable) Compliance
MINIMISE	o	+								
HIDE		+					o			
SEPARATE	o						o			
AGGREGATE	o	+								
INFORM				+	+				+	
CONTROL			o		+			+		
ENFORCE	+		+			+	+			o
DEMONSTRATE										+

Legend: +: covers principle to a large extent. o: covers principle to some extent.

Table 1. Mapping of strategies onto legal principles.

Design patterns Design patterns that implement this strategy are, for example, *privacy management systems* [20], and the use of logging and auditing.

6 Conclusions & Acknowledgements

We have derived eight privacy design strategies from an IT system perspective, taking the legal requirements as point of departure. The coverage of these legal principles by the design strategies is summarised in Table 1 (using the detailed description of each strategy in section 5).

As discussed before, not every legal data protection principle can be covered by a privacy design strategy, simply because the design of the system has no impact on that principle. Some data protection principles, like purpose limitation, are only partially covered by some of the strategies. Realising purpose limitation in full also requires procedural and organisational means.

With respect to design pattern coverage, we first observe that design patterns may belong to several design strategies. For example the *use pseudonyms* design pattern both implements the MINIMISE strategy and the HIDE strategy. In the course of our investigations we also observed huge differences between design strategies in terms of the number of design patterns known to implement them. For the strategies MINIMISE and HIDE, a large number of design patterns exist. This is not surprising, given the focus of most research in privacy enhancing technologies on these aspects of privacy protection. For the SEPARATE and CONTROL strategies on the other hand, no corresponding design patterns are known.

This paper discusses work in progress. In particular, further research will be performed to classify existing privacy design patterns into privacy design strategies,

and to describe these design patterns in more detail. Moreover, we have identified several implicitly defined design patterns (like *attribute based credentials*) that arise from our study of existing privacy enhancing technologies. Further developments and collaboration in this line of research will also be documented on our Wiki¹⁰. We would very much welcome contributions from the research community.

I would like to thank the members of the Privacy & Identity Lab¹¹ for discussions and valuable feedback. In particular I am grateful to Ronald Leenes, Martin Pekarek and Eleni Kosta for their detailed comments and recommendations that greatly improved this paper.

References

1. J. Borking. Der identity-protector. *Datenschutz und Datensicherheit*, 20(11):654–658, 1996.
2. Stefan Brands. *Rethinking Public Key Infrastructures and Digital Certificates; Building in Privacy*. MIT Press, 1 edition, 2000. ISBN 0-262-02491-8.
3. Frank Buschmann, Regine Meunier, Hans Rohnert, and Peter Sommerlad. *Pattern-Oriented Software Architecture, Volume 1: A System of Patterns*. John Wiley & Sons, 1996.
4. Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In Birgit Pfitzmann, editor, *EUROCRYPT*, volume 2045 of *Lecture Notes in Computer Science*, pages 93–118. Springer, 2001.
5. Ann Cavoukian. Privacy by design – the 7 foundational principles. Technical report, Information and Privacy Commissioner of Ontario, jan 2011. (revised version).
6. David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–88, 1981.
7. David Chaum, Amos Fiat, and Moni Naor. Untraceable electronic cash. In Shafi Goldwasser, editor, *CRYPTO*, LNCS 403, pages 319–327. Springer, 1988.
8. Communication COM (2007)228 from the Commission to the European Parliament and the Council. On Promoting Data Protection by Privacy Enhancing Technologies (PETs). (*Not published in the OJC*), 2007.
9. 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. *OJ C L*, 281:0031 – 0050, November 23 1995.
10. 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. *OJ C*, 102:24, April 5 2012.
11. Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1994.
12. Cornelia Graf, Peter Wolkerstorfer, Arjan Geven, and Manfred Tscheligi. A pattern collection for privacy enhancing technology. In *The 2nd Int. Conf. on Pervasive Patterns and Applications (PATTERNS 2010)*, Lisbon, Portugal, November 21–26 2010.
13. Seda Gürses, Carmela Troncoso, and Claudia Diaz. Engineering privacy by design. In *Conference on Computers, Privacy & Data Protection (CPDP 2011)*, 2011.

¹⁰ <http://wiki.science.ru.nl/privacy/>

¹¹ <http://www.pilab.nl>

14. Munawar Hafiz. A collection of privacy design patterns. In *Proceedings of the 2006 conference on Pattern languages of programs*, PLoP '06, pages 7:1–7:13, New York, NY, USA, 2006. ACM.
15. Munawar Hafiz. A pattern language for developing privacy enhancing technologies. *Softw. Pract. Exper.*, 2011. doi: 10.1002/spe.1131.
16. Jaap-Henk Hoepman. Privacy design strategies, October 2012. eprint arXiv:1210.6621. A preliminary version was presented at the Amsterdam Privacy Conference (APC 2012) and the Privacy Law Scholars Conference (PLSC 2013).
17. ISO/IEC 29100. Information technology – Security techniques – Privacy framework. Technical report, ISO JTC 1/SC 27.
18. B. Jacobs. Select before you collect. *Ars Aequi*, 54:1006–1009, December 2005.
19. P Kruchten. An ontology of architectural design decisions. In Jan Bosch, editor, *Proc. of the 2nd Groningen Workshop on Software Variability Management*, Groningen, The Netherlands, 2004.
20. Marco Casassa Mont and Siani Pearson. An adaptive privacy management system for data repositories. In Sokratis K. Katsikas, Javier Lopez, and Günther Pernul, editors, *Trust, Privacy and Security in Digital Business: Second International Conference, TrustBus 2005, Copenhagen, Denmark, August 22-26, 2005, Proceedings*, LNCS 3592, pages 236–245. Springer, 2005.
21. Organisation of Economic Co-Operation and Development. OECD guidelines on the protection of privacy and transborder flows of personal data, 1980.
22. Siani Pearson and Azzedine Benameur. Decision support for design for privacy: A system focused on privacy by policy. In *PrimeLife/IFIP Summer School 2010: Privacy and Identity Management for Life*, Helsingborg, Sweden, August 2010. (to appear).
23. Siani Pearson and Yun Shen. Context-aware privacy design pattern selection. In Sokratis K. Katsikas, Javier Lopez, and Miguel Soriano, editors, *Trust, Privacy and Security in Digital Business (TrustBus), 7th International Conference*, LNCS 6264, pages 69–80, Bilbao, Spain, August 30–31 2010. Springer.
24. Andreas Pfitzmann and Marit Hansen. Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management – a consolidated proposal for terminology (version v0.34 aug. 10, 2010). http://dud.inf.tu-dresden.de/Anon_Terminology.shtml.
25. D. J. Solove. A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3):477–564, 2006.
26. Sarah Spiekermann and Lorrie Faith Cranor. Engineering privacy. *IEEE Trans. Software Eng.*, 35(1):67–82, 2009.
27. Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):557–570, 2002.
28. US Federal Trade Commission. Privacy online: Fair information practices in the electronic marketplace, a report to congress, 2000.
29. G. W. van Blarckom, J. J. Borking, and P. Verhaar. PET. In G. W. van Blarckom, J. J. Borking, and J. G. E. Olk, editors, *Handbook of Privacy and Privacy-Enhancing Technologies - The case of Intelligent Software Agents*, chapter 3, pages 33–54. College bescherming persoonsgegevens, The Hague, The Netherlands, 2003.
30. Jeroen van Rest, Daniel Boonstra, Maarten Everts, Martin van Rijn, and Ron van Paassen. Designing privacy-by-design. Presented at the Annual Privacy Forum 2012, Limmasol, Cyprus.
31. A. Westin. *Privacy and Freedom*. Atheneum, New York, 1976.