Inter-temporal Privacy Metrics

Stefan Berthold

Faculty of Health, Science and Technology

Computer Science

DISSERTATION | Karlstad University Studies | 2014:63

Inter-temporal Privacy Metrics

Stefan Berthold

DISSERTATION | Karlstad University Studies | 2014:63

Inter-temporal Privacy Metrics

Stefan Berthold

DISSERTATION

Karlstad University Studies | 2014:63

urn:nbn:se:kau:diva-33972

ISSN 1403-8099

ISBN 978-91-7063-603-5

Distribution:
Karlstad University
Faculty of Health, Science and Technology
Department of Mathematics and Computer Science
SE-651 88 Karlstad, Sweden
+46 54 700 10 00

Print: Universitetstryckeriet, Karlstad 2014

WWW.KAU.SE

Inter-temporal Privacy Metrics

Stefan Berthold
Department of Mathematics and Computer Science, Karlstad University

Abstract

Informational privacy of individuals has significantly gained importance after information technology has become widely deployed. Data, once digitalised, can be copied, distributed, and long-term stored at negligible costs. This has dramatic consequences for individuals that leave traces in the form of personal data whenever they interact with information technology, for instance, computers and phones; or even when information technology is recording the personal data of aware or unaware individuals. The right of individuals for informational privacy, in particular to control the flow and use of their personal data, is easily undermined by those controlling the information technology.

The objective of this thesis is to study the measurement of informational privacy with a particular focus on scenarios where an individual discloses personal data to a second party, the data controller, which uses this data for re-identifying the individual within a set of other individuals, the population. We contribute with privacy metrics for several instances of this scenario in the publications included in this thesis, most notably one which adds a time dimension to the scenario for modelling the effects of the time passed between data disclosure and usage. The result is a new framework for inter-temporal privacy metrics.

The common dilemma of all privacy metrics is their dependence on the information available to the data controller. The same information may or may not be available to the individual and, as a consequence, the individual may be misguided in his decisions due to limited access to the data controller's information when using privacy metrics. The goal of this thesis is thus not only the specification of new privacy metrics, but also the contribution of ideas for mitigating this dilemma. However, an effective solution to the dilemma, if possible, would rather be a combination of technological, economical and legal means than a purely technical solution.

Keywords: privacy, unlinkability, metrics, uncertainty, valuation process, domain-specific language, anonymous communication.

Acknowledgements

This dissertation would not have been possible without the support of a number of people. First, I would like to thank my supervisor Simone Fischer-Hübner and my co-supervisor Stefan Lindskog for their advice and steady support in my research. Also, I am grateful for the collaboration with my co-authors, Rainer Böhme, Sebastian Clauß, Simone Fischer-Hübner, Reine Lundin, Leonardo Martucci, Stefan Köpsell, Tobias Pulls, and Ge Zhang.

Beside my supervisors and the co-authors, many colleagues added a lot to my inspiration, skills, and knowledge (not only) in privacy and security. Among these are Stefan Alfredsson, Julio Angulo, Marcel Cavalcanti de Castro, Peter Dely, Johan Garcia, Thomas Gloe, Hans Hedbom, Matthias Kirchner, Andreas Pfitzmann, Stefanie Pötzsch, Monika Sturm, and Antje Winkler, Philipp Winter, and Rose-Mharie Åhlfeld.

My research has been partly funded by the EU projects PRIME, FIDIS, PrimeLife, and A4Cloud; and by the PETweb II project of the Research Council of Norway. Also, I have received stipends for visiting the PETS symposium in Leuven, 2008, and for visiting the Summer School on Formal Methods for the Science of Security in Urbana, Illinois, 2013.

As my work was not done in social void, I owe my deepest gratitude to my family for their steady support during all my life and to my girlfriend for her love and patience during the last years. I am also lucky to have friends such as Matti and Frank Hilbert who introduced me to sport climbing and desperately try to get me on a trip to Lago di Garda since years, Sebastian Clauß who was my first regular climbing partner in Dresden, Lea Schneider, Stina Gustafsson, Toke Høiland-Jørgensen, Jonas Karlsson, and Philipp Winter who have become my regular climbing partners in Karlstad, Andreas Lavén who does not like climbing but qualifies as a reliable diving partner and is always good for a fresh interpretation of otherwise clear rules in board and card games, Jakob Krause who used to organise game evenings and invited me to a board game weekend in Dresden right after I moved to Karlstad (impossible to reject!), Alexis Crespel who after discovering the beauty of the Saxon Switzerland introduced me to the beauty of his Alsace, Peter and Nina Hjärtquist who became early and lasting friends when I moved to Sweden, Jens and Stephanie Lehmann as well as Jens Raschke who remained close frinds despite my move to Sweden, and Åsa Rangfeldt who has inspired and encouraged me to explore my skills in photography during the last years.

Karlstad, November 2014 Stefan Berthold

List of included publications

I. Stefan Berthold and Reine Lundin. Re-identification revisited. Under submission, 2014.

II. Stefan Berthold and Sebastian Clauß. Linkability Estimation Between Subjects and Message Contents Using Formal Concepts. In: Atsuhiro Goto (ed.), *Proceedings of the workshop on Digital Identity Management (DIM), co-located with the confernce on Computer and Communications Security (CCS)*, ACM, 2007. pp. 36–45

III. Stefan Berthold, Rainer Böhme, and Stefan Köpsell. Data Retention and Anonymity Services—Introducing a New Class of Realistic Adversary Models. In: Vacláv (Vašek) Matyáš, Simone Fischer Hübner, Dan Cvrček, and Petr Švenda (eds.), *The Future of Identity in the Information Society*, IFIP Advances in Information and Communication Technology, Springer, 2009. pp. 92–106

The publication won the best student paper award.

IV. Ge Zhang and Stefan Berthold. Hidden VoIP Calling Records from Networking Intermediaries. In: Georg Carle, Helmut Reiser, Gonzalo Camarillo, and Vijay K. Gurbani (eds.), *Proceedings of the Principles, Systems and Applications of IP Telecommunications (IPTComm)*, ACM, 2010. pp. 12–21

V. Stefan Berthold and Rainer Böhme. Valuating Privacy with Option Pricing Theory. In: Tyler Moore, David Pym, and Christos Ioannidis (eds.), *Economics of Information Security and Privacy*, chapter 10, Springer, 2010. pp. 187–209

The work has also been presented on the Workshop on Economics of Information Security (WEIS), 2009.

VI. Stefan Berthold. Towards a Formal Language for Privacy Options. In: Simone Fischer-Hübner, Penny Duquenoy, Marit Hansen, Ronald Leenes, and Ge Zhang (eds.), *Privacy and Identity Management for Life*, IFIP Advances in Information and Communication Technologies, Springer, 2011. pp. 27–40

VII. Stefan Berthold. The Privacy Option Language—Specification and Implementation. Karlstad University Studies, Report 2013:29, ISBN 978-91-7063-507-6, Department of Mathematics and Computer Science, Karlstad University, 2013.

Comments on my participation

Publication I The idea for this paper emerged from my discussions with Reine Lundin. I had the main ideas presented in the paper and discussed them further with Reine. Also, I am the principal author of the paper.

Publication II I am the principal author of all parts of the paper and received a number of useful comments from Sebastian Clauß.

Publication III This paper was a joint work of all three authors. I was the driving force behind the first ideas of this paper and later took the responsibility for adding the sections on the attacks and bringing the theoretical ideas together with the study in a coherent paper. Stefan Köpsell carried out the empirical study and drafted the results of it in German. He was also the main author of the sections about anonymity services in a nutshell and the legal background. Rainer Böhme contributed the section about probabilistic intersection attacks and supervised the study.

Publication IV The idea for this paper emerged from my discussions with Ge Zhang. We jointly elaborated a solution. While I contributed the ideas around the anonymity service, Ge added with his background in VoIP.

Publication V This paper was a joint effort with Rainer Böhme. I authored the main part of the paper, i. e., the related work on anonymity and unlinkability as well as Sections 3, 4, 5, and 6. Rainer authored the introduction, the related work on financial methods in information security, and the conclusions; and contributed his knowledge about economics to the main ideas of this paper.

Publication VI This paper was solely authored by myself.

Publication VII This report was solely authored by myself.

During the course of writing the publications, I received many constructive and motivating comments from my supervisors, other colleagues, and friends.

Other publications

- Stefan Berthold. 'Possibilistic Disclosure Attacks in Polynomial Time'. In: *Pre-proceedings of the FIDIS/IFIP Internet Security & Privacy Summer School*. Masaryk University, 2008
- Stefan Köpsell and Stefan Berthold. 'Public Key Infrastructure'. In: *The Future of Identity in the Information Society: Challenges and Opportunities*. Ed. by Kai Rannenberg, Denis Royer and André Deuker. FIDIS Summit Book. Springer, 2009, pp. 133–136
- Stefan Köpsell and Stefan Berthold. 'Electronic Signatures'. In: *The Future of Identity in the Information Society: Challenges and Opportunities*. Ed. by Kai Rannenberg, Denis Royer and André Deuker. FIDIS Summit Book. Springer, 2009, pp. 136–138

- Stefan Berthold and Sebastian Clauß. 'Privacy Measurement'. In: *Digital Privacy, PRIME – Privacy and Identity Management for Europe*. Ed. by Jan Camenisch, Ronald Leenes and Dieter Sommer. Vol. 6545. LNCS. Springer, 2011
- Stefan Berthold. *Towards Inter-temporal Privacy Metrics*. Licentiate thesis 2011:25. Karlstad University Studies, ISBN 978-91-7063-357-7. Karlstad University, 2011
- Simone Fischer-Hübner and Stefan Berthold. 'Privacy-Enhancing Technologies'. In: *Computer and Information Security Handbook*. Ed. by John R. Vacca. 2nd edition. Morgan Kaufmann, 2012. Chap. 43, pp. 755–772
- Stefan Berthold, Simone Fischer-Hübner, Leonardo Martucci and Tobias Pulls. *Crime and Punishment in the Cloud: Accountability, Transparency, and Privacy*. Proceedings of the International Workshop on Trustworthiness, Accountability and Forensics in the Cloud (TAFC). 2013

Contents

Introductory Summary

1 Introduction

Determining when people first thought about their privacy is probably impossible, but it is reasonable to assume that people ever since sought to assess which actions lead to better privacy and which to worse. The ways of measuring privacy depend on the category of privacy, e. g., physical privacy or informational privacy. In this collection thesis, we focus on the latter, informational privacy.

Informational privacy of individuals is in technical contexts understood as the control of the individuals over their personal data, including the disclosure, flow, and storage. As information technology becomes widely deployed, retaining this control gains significant importance for individuals. In information technology, data, once digitalised, can be copied, distributed, stored, and filtered at negligible costs. This becomes problem for the individuals' privacy when the individuals loose control over their data. Privacy metrics can help individuals to evaluate the consequences of disclosing personal data and thus help them to make informed decisions about the personal data they want to disclose.

In the included articles of this thesis, we study the measurement of informational privacy. We focus on scenarios where the individual discloses personal data to a second party, the data controller, which uses this data for re-identifying the individual within a set of other individuals, the population. We contribute with new privacy metrics and the corresponding attacker models. These metrics can be used as transparency and data minimisation tools. Also, we contribute with a novel framework for inter-temporal privacy metrics, in which existing privacy metrics can be integrated, and define a privacy-preserving contract language for data disclosures which can be used to specify the inputs for inter-temporal privacy metrics.

The remainder of the introductory summary is structured as follows. Section 2 (Background and related work) introduces the basic terminology and concepts we use and refer to in this thesis and discusses related work. Section 3 (Research questions) defines the research questions addressed in this thesis, followed by a discussion in Section 4 (Research methods) of research methods that we applied for addressing these challenges. Section 5 (Contributions) states the contributions to the research field of privacy metrics. Section 6 (Summary of the publications) summarises the included publications. Section 7 (Conclusions and future work) provides conclusions and outlines future work.

2 Background and related work

This section summarises the basic terms and concepts used in this thesis and refers to related work.

2.1 Informational privacy

Since the central concept of this thesis is *informational privacy*, we start the background section with defining it. Though there is a large body of work using the term, there is no consolidated definition of privacy. One of the earliest definitions in modern scientific literature was given by Westin who has defined privacy as

> "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others." [45]

The German Federal Constitutional Court has defined privacy as the right of *informational self-determination*:

> "[The] constitutional right of informational self-determination guarantees the permission of an individual to determine on principle himself on the disclosure and use of his personal information." [21]

These notions have been widely adopted and we will thus not redefine informational privacy, but rather focus on the facets of it which are discussed in the publications included in this thesis.

2.2 Anonymity

Pfitzmann and Hansen [32] were maintaining a terminology overview for research in the informational privacy domain. The continuous updates of this terminology have been widely adopted in other scientific work since its first publication in 2000 [33]. Pfitzmann and Hansen define anonymity as follows:

> "Anonymity of a subject means that the subject is not identifiable within a set of subjects, the anonymity set." [32]

The subjects in this definition are the actors in a network that either send or receive messages. If, for instance, the sending subject of a message cannot be determined among a set of other sending subjects, then the sending subject of this message is anonymous and the set of sending subjects in which the sender of the message is hiding is the *anonymity set*.

The size of the anonymity set is a measurement of anonymity. For sizes greater than 1, the sending subject is considered anonymous. An anonymity set size of exactly 1 means that the sending subject is *identified*. Though not explicitly mentioned in [32], the anonymity set notion was most likely motivated by the way anonymity was established in Chaum's DC network [8]. In the DC network, all participants are equally likely the sender of each message. Neither insiders nor outsiders can obtain more information from the network for narrowing down the sender. Chaum pioneered the research on privacy-enhancing technology (PET). Besides the DC network, he also laid the foundations for another PET, the mix [9]. Mixes are network proxies with

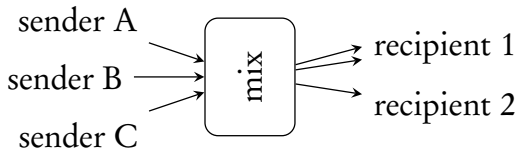

Figure 1: Illustration of a mix. The mix operation makes it impossible to trace an incoming message through the mix and relate it to an outgoing message.

some interesting properties, see [18, 16] for surveys of different types of mixes and Figure 1 for an illustration of a mix. The term 'mix' refers to the mode of operation of these proxies whose main purpose it is to disturb (or 'mix') the observable relation between incoming and outgoing messages. A user who is sending a message to the mix can thus be certain (to the extend security is provided by the methods used, e. g., public-key cryptography [38]) that his message can not only be related to him but to the same extent to all other users who sent a message to the mix at the same time. The user is thus hidden in the anonymity set of all senders.

Mix users are not required to send messages continuously over time, i. e., over several mix rounds. This allows a number of attacks and though these attacks do not help revealing a sender–message relation, they are efficient in uncovering sender–recipient relations. The simplest attack is intersecting two or more anonymity sets which have been observed when a specific recipient got a message. This attack and variations of it are known as long-term intersection attacks [6]. If the accuracy of the result is only relevant up to a certain error margin, there is an even more efficient class of attacks, the statistical disclosure attacks [15, 27] which also take the probabilities into account of mix users in each anonymity set being the sender of a message to a specific recipient. In the mix round illustrated in Figure 1, we see that all senders submitted one message, but recipient 1 received two message while recipients 2 got only one. The probability of sender A being the sender of a message to recipient 1 is thus $\frac{2}{3}$ while the probability of A being the sender of a message to recipient 2 is only $\frac{1}{3}$. This probability distribution is not represented in the anonymity set notion.

A more general notion is the degree of anonymity [40, 19] which uses entropy in Shannon's [41] sense. The entropy of the probability distribution of individuals within the anonymity set can be interpreted as the expected information an attacker would learn when a relation between sender and recipient is revealed to him. While the attacker has not discovered this relation, the entropy can be understood as a measure of the expected effort a successful attack on the mix would require. A third interpretation of the entropy is the (logarithm of the) expected anonymity set size provided by the mix round.

These interpretations and other entropy-metrics, such as cross entropy, the Kullback-Leibler divergence, and combinations thereof, are further discussed in Publication I. The anonymity set notion is reused in Publication II and extended with a clustering for data disclosures. Also, the effect of long-term intersection attacks on the sizes of the anonymity sets is analysed in Public-

ation III. In Publication IV, we use the anonymity set notion to argue for protocol-specific low-latency mixes. And in Publication V, we extend entropy metrics with a time dimension.

2.3 Unlinkability

Pfitzmann and Hansen define unlinkability as follows:

> "Unlinkability of two or more items of interest (IOIs, e. g., subjects, messages, actions, ...) from an attacker's perspective means that within the system (comprising these and possibly other items), the attacker cannot sufficiently distinguish whether these IOIs are related or not." [32]

Unlinkability generalises anonymity, i. e., unlinkability describes the non-existence of a relation between two items of interest (possibly including subjects) and anonymity describes the non-existence of a relation between a subject and an item of interest (possibly another subject). Thus, every attack on anonymity is at the same time an attack on unlinkability, but anonymity metrics must be seen as special cases of unlinkability metrics. For instance, Steinbrecher and Köpsell [42] use equivalence class membership relations to model the relations between items of interest. They use probability distributions on these membership relations to model the degree of unlinkability.

The generalisation from subjects and messages to items of interest allows to abstract away from pure network-based scenarios and more generally focus on the linkability of personal data to individuals instead. Early work motivated by this scenario lead to anonymity metrics for statistical databases [39, 17, 21, 43] and has been adapted [11] to the unlinkability notion later. This new scenario was also the motivation for research on a branch of identity management systems that support individuals in managing their personal data disclosure with regard to anonymity and unlinkability. The first solutions for such privacy-enhancing identity management [7, 12] lead to a refinement of the requirements [23] and culminated in several projects [22, 26, 31] funded by the European Community and other institutions [30].

Even though the objective of Publication II is proposing anonymity metrics, it uses several ('formal scaling') steps where no subject is involved. These intermediate steps can be understood as unlinkability metrics. In Publication I and V, we use anonymity metrics as an example which can easily be generalised to unlinkability metrics.

2.4 Information asymmetry and transparency

The purpose of PETs, such as DC networks and mixes, is to minimise the flow of personal data[1] on the technical level. The term 'data minimisation'

[1] Broader definitions of PETs can be found in the literature that also capture technology such as access control, identity management, and even TETs (described later). In this thesis, we restrict the meaning of PETs to data minimisation.

coined in the privacy community, derived from the proportionality principle of modern constitutional states, and used in legislations such as the German 'Bundesdatenschutzgesetz' refers to the strategy of individuals choosing the smallest set of personal data to disclose to a data controller for achieving a specific goal, e. g., medical treatment, selling or purchasing goods of any kind or subscribing and using a service.

The effects of data minimisation can be better explained in the context of contract theory [3]: in a market with asymmetric distribution of information among the parties, the ones with more information have an incentive to use their information against the parties with less information [1]. A data controller sharing personal data of an individual with a third party without informing the individual creates an information asymmetry in which the controller possesses more information than the individual, since only the data controller would know about the new distribution of the data. The individual (benefits or) suffers from the consequences of the distribution of the data, i. e., the data controller creates (positive or) negative externalities [20] for the individual [25, 44]. PETs help to reduce information asymmetries, and thus negative externalities, by reducing the flow of data from the individual to the data controller.

This view suggests that data minimisation is not the only strategy for reducing information asymmetries, but has a sibling with equivalent effects on data asymmetries: increasing transparency, i. e., increasing the flow of information from the data controller to the individual. Bellotti and Sellen [5] discuss both strategies and the relation between them.

While PETs are effective at the time of the data disclosure, transparency-enhancing tools (TETs) may be effective either as predictive tools (ex ante) before the data disclosure or as retrospective tools after the fact (ex post) [24]. Predictive TETs can be used by individuals as decision support tools that anticipate the future consequences of a data disclosure while retrospective TETs inform the individuals about the de facto consequences of their data disclosures. Both types can be mixed so that consequences of past data disclosures can be taken into account for future data disclosures.

The metrics in Publication I, II and V can be conceived as predictive TETs as well as the outcome of the experiment in Publication III and the contract language in Publication VI.

2.5 Privacy policy languages

Privacy policies describe which personal data is collected by a data controller, for what purpose, and what is going to happen with the data. The policies can be seen as predictive TETs which inform the individual about the intended usage of the data by the data controller. With this information, the individual can choose among data controllers by choosing the privacy policy that matches the individual's preferences best. The individuals' choice of a data controller could even be automated, e. g., as part of an identity management system, if preferences and policies are machine-processible.

Several specifications of formal languages for privacy policies and preferences have been proposed. Among the first privacy policy languages and the first that became a standard was the platform for privacy preferences (P3P) [37, 14]. It defines a number of assertions, partly formalised and partly free text assertions, for the content category of the collected data, purpose of the data collection and the processing, whether the data is used in a personally identifiable way, and which parties receive the data, among others [37]. In conjunction with P3P, 'a P3P preference exchange language' (APPEL) [13] has been proposed as the corresponding privacy preference language. The current working draft of the language allows to specify the P3P assertions acceptable for the individual by stating simple expressions.

Another privacy policy language is the enterprise privacy authorization language (EPAL) [4] which contributed a new concept to privacy policies, the obligation. According to the specification, obligations are additional steps to be taken when a certain action is performed [4].

EPAL has been discontinued and remained a submission for standardisation. A competing language for privacy policies, the extensible access control markup language (XACML) [28, 29], has become a standard earlier and has been found more flexible than EPAL [2].

In contrast to P3P, both EPAL and XACML are access control languages. This line has been continued in the PrimeLife policy language (PPL) [36] and its extension with provisions for accountability (A-PPL) [10] which extend XACML and specify a number of subdialects, including an equivalent for privacy preferences. The access control rules in all four languages are optimised for efficient decisions on access rights, thereby answering the question what data *may be used*, but are considerably less suited for the anticipation of access requests, i. e., the anticipation of what data is requested, how often, and when, and thereby not answering the question what data is *will actually be used*. The languages' use as a predictive TET is therefore limited.

In Publication VI and VII, we propose a language for privacy contracts which is better suited as a predictive TET and provides a more general model for obligations than EPAL and XACML (including (A-)PPL).

3 Research questions

This thesis addresses the following research questions.

1. *How can we reduce information asymmetries between data subjects and data controllers?*

   This question aims at finding new designs or design patterns for information technology that decrease the information gap which may occur when a data controller processes or forwards personal data of an individual. The information gap often leads to advantageous situations for the data controller in which it possesses more information than the individual (knows) about (or derived from) the personal data. Possible solutions to this question are discussed in all seven publications, in the

form of transparency-enhancing technologies and data minimisation tools. Publication I, III, V, VI, and VII provide predictive TETs and Publication II outlines a solution which can work as predictive as well as a retrospective TET. In addition, the canonical form of privacy option contracts in Publication VI and VII, and the mix design in Publication IV provide tools for data minimisation, thus as PETs.

2. *How can re-identification risks be assessed depending on the attacker model?*

   This question aims at finding reasonable definitions of privacy metrics that either provide qualitative or quantitative assessments of the risk of being re-identified by personal data. This question has been addressed in Publication I with entropy-based metrics and extended in Publication V. Publications VI and VII specify a prototype in which these metrics can be deployed in practice. Publication II addresses this research question with a anonymity set-based metric. Publications IV and III address this research question by providing qualitative statements about the re-identification risk in practical (Publication III) and hypothetical (Publication IV) anonymity networks.

3. *What are the effects of time on the validity of personal data and how can these effects be captured in privacy metrics?*

   This question deals with the knowledge of data controllers that obtain personal information for later usage. Provided that some personal data is volatile, the data controller has to account for data that, when used, is not applying to the data subject anymore. This has consequences for the privacy of the data subject over time, the data subject's *inter-temporal privacy*. In Publication II and III, we assume that privacy once lost cannot be regained. In reality, however, the longer data is stored without validation or update, the greater is the likelihood that the stored data is invalid. For the data controller, this means that the data becomes increasingly useless by time and the data subject thus implicitly regains privacy. A mathematical model of this process gives interesting insights into how the expected usefulness of the data for the data controller and the expected infringement on the privacy of the data subject can evolve over time. We discuss an approach of such a mathematical model in Publication V. In Publication VI and VII, we also discuss the effects of time in privacy contracts. The contracts described in these publications are reduced to canonical form depending on the operators of the contract language, including a number of time operators.

4. *What are the primitives of a contract language for data disclosures?*

   The market of personal data is mostly governed by regulation and policies on service or company level, but bilateral and specific contracts between data controllers and data subjects cannot be found as often. However, both the data controller and the data subject could benefit from specific contracts, if these are human- and machine-readable. The

data controller would benefit from accurately valuating the database of personal data including all effects that time is expected to have. The data subject benefits from learning not only abstractly what could possibly be done with the data, but more specifically what is planned to be done with the data. This research question is addressed in Publications VI and VII.

4 Research methods

The main research method used in this thesis is the *mathematical* method [34]. Besides that, we used the *scientific* method [35]. Both methods have a lot in common, but differ in their details. We briefly discuss them in this section.

Both methods consist of four steps. The first step in the mathematical method is *understanding* the problem. It is similar to the first step, *characterisation* of the problem, in the scientific method. In both methods, this step is to identify the *unknown*, the data and conditions, and to formulate them in suitable notation.

The second step in the mathematical method is devising a *plan* (also referred to as *analysis*), and in the scientific method, it is the formulation of the *hypothesis*. In both cases, the second step can be summarised as finding a connection between the data and the unknown. While the hypothesis in the scientific method usually is a mathematical model, the plan in the mathematical method can also be an analogy to a similar (or related) problem or a proof idea.

The third step in both methods can be different. In the mathematical method, it is *carrying out* the plan (synthesis), i. e., specifying the analogy or carrying out the proof. It can also be a *deduction* as the third step in the scientific method is, i. e., instantiating the hypothesis to predict the outcome of laboratory experiments or observations in the nature.

The forth step in the mathematical method is *looking back*. Its main purpose is to review the results from the previous steps, i. e., validating or verifying the results, possibly simplifying the plan, and look for other applications of the same solution, i. e., generalising the solution. In the scientific method, the fourth step is mainly concerned with *validating* the predictions from the previous step in experiments.

In Publications I, II, IV, and V–VII we applied the mathematical method. In Publication III, we applied the scientific method. In Publication I, each section of the main matter (Sections 3–6) runs through the cycle of the mathematical method by defining an attacker model (understanding the problem), matching it with entropy-based metrics from information theory (analysis), discussing the bounds, evaluating the metrics, and comparing them to the other attacker models (synthesis). Also, the publication's Section 7 reflects on the general applicability of all proposed metrics (synthesis and review), and Section 8 reviews the proposed metrics in the context of other privacy metrics (review).

In Publication II, Section 1 motivates a privacy metric that can cope with content data (understanding the problem), for instance data from profiling and counter-profiling, and proposes an approach using a particular cluster method (analysis). The Sections 2–4 instantiate the cluster method in various ways (synthesis) and demonstrate the results visually as diagrams as well as formally in set notation (review).

In Publication III, we use the scientific method to carry out an experiment with data gathered from an operational anonymity service. The problem statement, a back then newly introduced law that made it mandatory to retain (internet) traffic data at the service provider and provide it upon request to law-enforcement, is described in Section 1 (problem characterisation). The hypothesis, negative impact on the users' privacy through well-known attacks, is introduced in Section 1 (hypothesis). The law, anonymity services, and the attacks are described together with their expected consequences in the Sections 2–5 (instantiation). The consequences are validated by an experiment for obtaining realistic traffic meta data and by a simulation of the attacks in Sections 6 and 7 (evaluation).

In Publication IV, we use the mathematical method and frame the problem by describing a communication model in Section 2 and attacker models in Section 3 (understanding the problem). In Section 4, we describe ways to mitigate the problem (analysis) and instantiate these ideas with an example solution (synthesis). In Sections 5 and 6, we discuss the scope of our solution and place it between other work in the field (review).

In Publication V, we develop in Section 3 the analogy between data disclosure and financial option contracts by briefly discussing the nature of data disclosures (understanding the problem), revealing the similarity to financial options (analysis), and matching the concepts that describe financial options to the terms that describe data disclosures (synthesis). The analogy is then used as the anchor (understanding the problem) for matching and using the valuation methods from finances with privacy metrics. This is done and evaluated in examples in Sections 4 and 5 (analysis and synthesis). In the Sections 6 and 7, the results are discussed and subsequent problems are identified (review).

In Publication VI, we look at the problem of automated privacy metrics evaluation in the context of complex data disclosures (Section 1, understanding the problem) and build on top of the analogy of Publication V to adapt a language for financial contracts for data disclosure contracts (Section 2, analysis and synthesis). The Sections 3–6 are showing or proving (Section 5) properties of the language (review). Each of these sections is also following the mathematical method.

In Publication VII, we use the language outline from Publication VI as the blueprint (understanding the problem and analysis) for a proof of concept implementation of the language framework in the Parts I and II (synthesis) and an instance of this framework in Part III (synthesis and review).

Also, the thesis as a whole is following the mathematical method where Publications I–IV can be seen as the analysis step, Publication V and VI as the synthesis and Publication VII as the review step.

5 Contributions

Transparency and data minimisation. We provide transparency tools in three categories (contributing to research question 1), privacy metrics which can be deployed as both predictive and retrospective transparency tools, analysis of anonymity networks which can be understood as predictive transparency tools, and a privacy contract language which also are predictive transparency tools. The privacy metrics and the analysis of anonymity networks (Publications I–V) assess the risk to be re-identified. The contract language (Publications VI and VII) can be used to establish clear agreements about when and what data is used for which purpose.

All predictive transparency tools can also be used as a means for minimising the disclosed data and information. Moreover, technology that attempts to address the problems of other technology creates new problems. In terms of data minimisation, this is reflected in paradoxical situations when data minimisation tools create new information leaks. The contracts represented in our contract language are meta data about the disclosed data. In order to minimise the information that can be derived from the meta data, we provide means to reduce the contracts to their canonical form, thereby effectively stripping off any information about the origin or the process of creating the contract.

Privacy metrics and attacker models. We define new realistic attacker models and privacy metrics for these attacker models, contributing to research question 2. The privacy metrics range from set-based metrics (Publications II, III, and IV) to entropy-based metrics (Publications I and V). The attacker models range from abstract attacker models (Publication I and V) to attacker models that are specific to realistic anonymity systems (Publication III) and hypothetical ones (Publication IV). All attacker models are matched with corresponding privacy metrics. Since all set-based metrics can be generalised to entropy-based metrics (with the uniform probability distribution), the set-based metrics can be understood as a base that is extended and generalised by our entropy-based metrics. Moreover, the metrics without clear inter-temporal view (Publication I, II, and IV) are generalised in the metrics with deal more (Publication V) or less (Publication IV) explicit with time within attacks.

Inter-temporal privacy model. We provide a new view on privacy metrics by explicitly accounting for the time passed between data disclosure and data usage, contributing to research questions 2 and 3. This leads to an interesting analogy between (deliberate) data disclosures and option contracts from finances. In both cases, the value, i. e., the result of a privacy measurement and the monetary value of the financial contract respectively, is determined by the future development of the environment in which the data will be used and the contract will be executed. We found that a hole class of valuation methods can be mapped from the finance domain to the privacy domain, thereby providing access to a new set of instruments that has not been explored

before in privacy metrics research. The analogy and the resulting model are presented in Publication V. Despite we used Shannon entropy for the evaluation of the model, it is not restricted to any specific entropy-based metric, but can be evaluated with any metric that fits the attacker model. In this regard, Publication V provides a framework of metrics and Publications I–IV constitute more (Publication I) or less (Publications II–IV) intuitive examples of basic metrics that can be hooked in this framework.

Privacy contract language. We define and evaluate a privacy contract language which captures inter-temporal aspects of data disclosures, contributing to research questions 3 and 4. The language can be understood as both a stricter form of privacy policies and the specification of the input format for the inter-temporal privacy model. In contrast to privacy policy languages that provide the access control instructions for the handling and processing of personal data, privacy contracts *are* the instructions for privacy-related data handling and processing. In the details, privacy contracts turn some concepts of privacy policies upside-down, for instance rights and obligations: privacy policies define rights with attached obligations, in contrast to privacy contracts which are obligations unless accompanied with the option to abandon the contract, then the contract would become a right. The recursive structure of contract modification functions like the option to abandon contracts provides straight forward ways to evaluate the contract with privacy metrics as presented in Publication V or execute the contract. Also, the contract language design makes it easy to normalise equivalent contracts to their canonical form and thereby avoiding to create new covert channels in data disclosures by using the contract language. The contract language is specified and implemented in Publications VI and VII and the contracts can serve as input for any of the metrics defined in Publications I–V.

6 Summary of the publications

In this section, the included publications and the relations among them are briefly summarised.

Publication I – Re-identification revisited

In this paper, we present an introduction to entropy-based metrics for re-identification. We analyse when these metrics can and should be applied and when the metrics over- or underestimate the re-identfication workload. Re-identification metrics are directly or indirectly used in all included publications in this thesis, either for measuring linkability on the application layer or to measure anonymity on the link layer.

Publication II – Linkability Estimation Between Subjects and Message Contents Using Formal Concepts

In this paper, we present an approach to linkability of individuals' data disclosures. The approach uses graph threory to produce a kind of anonymity set notion. Also, the mechanisms used provide an ontological clustering of the data disclosures where similar disclosures are close and unlike disclosures are distant nodes in the graph. This extension of anonymity sets with similarity clustering provides information that can be used for determining probabilities in entropy-based approaches. As such, the work in this paper can also be seen as a step towards entropy-based metrics.

Publication III – Data Retention and Anonymity Services

In this paper, we present empirical results on the anonymity that can be provided within the terms of the Germany's first implementation of the European Data Retention Directive (2006/24/EC). This implementation required that, e. g., internet service providers retain traffic data and thus specified a model for data collection. It also specified an attacker model by regulating the conditions under which the data could be accessed and by whom. The study is performed using an operative anonymity service and real user data. The degree of anonymity has been examined under the attacker model specified in the implementation of the European Data Retention Directive and under less restrictive assumptions about the attacker.

Publication IV – Hidden VoIP Calling Records from Networking Intermediaries

In this paper, we discuss how VoIP traffic can be anonymised. VoIP produces a lot and in specific conditions evenly shaped traffic, thus eliminating characteristic problems we have when anonymising web or arbitrary low-latency traffic on the link layer. This makes VoIP a promising target for anonymising the data by means of mixes. It turns out that the start and the end of VoIP sessions are most sensitive with regard to anonymity. VoIP is an example which demonstrates that anonymisation methods for specific protocols may outperform general methods in two ways: first, they can utilise the specifics of the native protocols whenever beneficial, and second, they can cope with the sensitive parts of the specific protocol.

Publication V – Valuating Privacy with Option Pricing Theory

In this paper, we present a framework for measuring informational privacy. The actual measurement of informational privacy can be done by any reasonable metric that deals with attribute frequencies. The framework adds the time dimension, i. e., a model for uncertainty about the individual's attributes induced by the development of the individual and the population of all individuals over time. Our uncertainty model builds on the analogy to uncertainty

models in option pricing known from economics. The framework is thus called Privacy Option.

Publication VI – Towards a Formal Language for Privacy Options

In this paper, we present a formally specified language, POL, for Privacy Options. It particularly focuses on the aspects of time in Privacy Options. It is similar to privacy policy languages such as P3P and PrimeLife's PPL as both types of languages specify the rights and obligatons of data controllers, but differs in its design as a contract model rather than an access control model. In POL, rights are defined as obligation contracts with the freedom to nullify them, i. e., rights and obligations are expressed with the same vocabulary, which allows defining obligations in a more flexible way than known from privacy policy languages. Besides, we present a canonical form of POL contracts which can be used to normalise the syntax of contracts and thus avoid unintentionally creating covert channels which could create new privacy risks otherwise.

Publication VII – The Privacy Option Language

In this report, we present our implementation of the Privacy Option Language as an embedded domain-specific language. It is used as a proof of concept demonstrator of the language's characteristics, e. g., the canonical form, and its semantics. Also, the reduction rules that reduce contracts to canonical form have been refined compared to Publication VI. For the simulation of simple identity management, simPOL, a simple POL instance, implements a playground for testers and developers.

7 Conclusions and future work

In this thesis, we have shown that privacy metrics are an important research area with challenging research questions. Also, we have presented several novel approaches for privacy metrics in Publication I, II, and V. An important contribution was to underpin the uncertainty about the validity of collected data induced by time with a mathematical model, the Privacy Option. Various parameters of this model can be specified in terms of contracts in the Privacy Option Language and privacy measurement can be defined as semantics of these contracts. Other contract semantics are conceivable, e. g., the data management semantics defined in Publication VI and VII. We are convinced that future contract languages for privacy-enhancing identity management have to provide and cope with these two semantics as well.

In Publication III, we have measured privacy with empirical data and with regard to a realistic (law-abiding) attacker model and an operational anonymity service. The conclusion is that, despite the fears in public, privacy can be preserved quite well against the law-abiding attacker. However, attacks on the

same data without the restrictions of the law-abiding attacker may have more severe consequences in terms of privacy decrease.

In Publication IV, we analysed protocol-specific anonymity and whether it can be implemented by mixes in the case of Voice over IP (VoIP). It is interesting to see that VoIP, though demanding in terms of bandwidth and latency, matches very well to mix settings. In fact, the requirements in bandwidth make it possible to deploy low-latency mixes with random packet shuffling, a feature that is not supported by most operative protocol-independent low-latency mix networks. Indeed, this is more a primer for protocol-aware mix networks than for several protocol-specific ones. The advantage of protocol-awareness is that packages from different protocols can be mixed, and therefore the critical mass of users can easier be maintained, and at the same time the packages can be routed through the net with their optimal strategy. However, protocol-aware mixes only make sense when the abuse of the protocol-dependent packet routing bias is punished, either in terms of detection and for instance exclusion from the network or by a diminished anonymity for the abuser. Protocol-aware mixes will therefore be an interesting topic for future research.

Publication II extends the anonymity set notion with clustering methods. One of the most interesting challenges is to generalise the methods such that imperfect knowledge can be modelled in all relations, i. e., imperfect knowledge about deducible data items as well as imperfect knowledge about the linkability of data to an individual. Publication V can be seen as a follow-up in this regard, since modelling imperfect knowledge is the main subject of this publication. In order to present the metrics in Publication V in a compact way, we limited our model of personal data to the most basic case, i. e., one attribute with two attribute values. In future work, this will be generalised to an arbitrary number of attributes and attribute values.

Future work for the contract language specified in Publication VI may focus on new languages that wrap the contracts and put them in a context, e. g., specifying the contract parties. New languages that can be wrapped by the contract language are also conceivable, e. g., expression languages that refine the data model so that contracts can be made that allow, disallow, or enforce specific operations on the data.

References

[1] George Arthur Akerlof. 'The Market for "Lemons": Quality Uncertainty and the Market Mechanism'. In: *The Quarterly Journal of Economics* 84.3 (1970), pp. 488–500.

[2] Anne Anderson. *A Comparison of Two Privacy Policy Languages: EPAL and XACML*. Technical Report. Sun Microsystems Laboratories, 2005.

[3] Kenneth Joseph Arrow. 'Uncertainty and the Welfare Economics of Medical Care'. In: *The American Economic Review* LIII.5 (1963), pp. 941–973.

[4] Paul Ashley, Satoshi Hada, Günter Karjoth, Calvin Powers and Matthias Schunter. *Enterprise Privacy Authorization Language (EPAL 1.2)*. Member Submission. W3C, 2003.

[5] Victoria Bellotti and Abigail Sellen. 'Design for privacy in ubiquitous computing environments'. In: *Proceedings of the third conference on European Conference on Computer-Supported Cooperative Work (ECSCW'93)*. Ed. by Giorgio de Michelis, Carla Simone and Kjeld Schmidt. Kluwer Academic Publishers, 1993, pp. 77–92.

[6] Oliver Berthold. 'Effiziente Realisierung von Dummy Traffic zur Gewährleistung von Unbeobachtbarkeit im Internet [An efficient implementation of dummy traffic to ensure unobservability on the Internet]'. Diploma thesis. Technische Universität Dresden, Faculty of Computer Science, Institute for Theoretical Computer Science, 1999.

[7] Oliver Berthold and Hannes Federrath. 'Identitätsmanagement'. In: *E-Privacy: Datenschutz im Internet (DuD-Fachbeiträge)* (2000), pp. 189–204.

[8] David Chaum. 'The dining cryptographers problem: Unconditional sender and recipient untraceability'. In: *Journal of Cryptography* 1.1 (1988), pp. 65–75.

[9] David Chaum. 'Untraceable electronic mail, return addresses, and digital pseudonyms'. In: *Communications of the ACM* 24.2 (1981), pp. 84–90.

[10] Roman-Alexandre Cherrueau et al. *Policy Representation Framework*. A4Cloud Project Deliverable D34.1. 2013.

[11] Sebastian Clauß. 'A Framework for Quantification of Linkability Within a Privacy-Enhancing Identity Management System'. In: *Emerging Trends in Information and Communication Security (ETRICS)*. Ed. by Günter Müller. Vol. 3995. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2006, pp. 191–205.

[12] Sebastian Clauß and Marit Köhntopp. 'Identity Management and its Support of Multilateral Security'. In: *Computer Networks* 37.2 (2001), pp. 205–219.

[13] Lorrie Faith Cranor, Marc Langheinrich and Massimo Marchiori. *A P3P Preference Exchange Language 1.0 (APPEL1.0)*. Working Draft. W3C, 2002.

[14] Lorrie Faith Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall and Joseph Reagle. *The Platform for Privacy Preferences 1.0 (P3P1.0) Specification*. Recommendation. W3C, 2002.

[15] George Danezis. 'Statistical disclosure attacks: Traffic confirmation in open environments'. In: *Proceedings of Security and Privacy in the Age of Uncertainty (SEC2003)*. Ed. by Stefano Gritzalis, Sabrina De Capitani di Vimercati, Pierangela Samarati and Georgios Katsikas. Kluwer, 2003.

[16] George Danezis and Claudia Díaz. *A Survey of Anonymous Communication Channels*. Tech. rep. MSR-TR-2008-35. Microsoft Research, 2008.

[17] Dorothy Elizabeth Denning, Peter James Denning and Mayer Dlugach Schwartz. 'The Tracker: A Threat to Statistical Database Security'. In: *ACM Transactions on Database Systems* 4.1 (1979), pp. 76–96.

[18] Claudia Díaz and Bart Preneel. 'Taxonomy of Mixes and Dummy Traffic'. In: *Proceedings of I-NetSec04: 3rd Working Conference on Privacy and Anonymity in Networked and Distributed Systems*. Springer, 2004.

[19] Claudia Díaz, Stefaan Seys, Joris Claessens and Bart Preneel. 'Towards measuring anonymity'. In: *Workshop on Privacy Enhancing Technologies*. Ed. by Paul Syverson and Roger Dingledine. Vol. 2482. LNCS. Springer, 2002.

[20] Howard Sylvester Ellis and William Fellner. 'External Economies and Diseconomies'. In: *The American Economic Review* 33.3 (1943), pp. 493–511.

[21] Simone Fischer-Hübner. *IT-security and privacy: Design and use of privacy-enhancing security mechanisms*. Vol. 1958. LNCS. Berlin Heidelberg: Springer, 2001.

[22] Marit Hansen and Henry Krasemann. *Privacy and Identity Management for Europe – PRIME White Paper*. Project Deliverable D15.1.d. Privacy and Identity Management for Europe (PRIME, Project within the European Community's 6th Framework Program, No. 507591), 2005.

[23] Marit Hansen et al. 'Privacy-Enhancing Identity Management'. In: *Information Security Technical Report* 9.1 (2004), pp. 35–44.

[24] Mireille Hildebrandt. *Behavioural Biometric Profiling and Transparency Enhancing Tools*. Project Deliverable 7.12. Future of Identity in the Information Science (FIDIS, Network of Excellence within the European Community's 6th Framework Program, No. 507512), 2009.

[25] Xiaodong Jiang, Jason Hong and James Landay. 'Approximate Information Flows: Socially-Based Modeling of Privacy in Ubiquitous Computing'. In: *UbiComp 2002: Ubiquitous Computing*. Ed. by Gaetano Borriello and Lars Holmquist. Vol. 2498. LNCS. Springer, 2002, pp. 176–193.

[26] Eleni Kosta and Jos Dumortier. *Contextual Framework*. Project Deliverable D2.3. Privacy and Identity Management for Community Services (PICOS, Project within the European Community's 7th Framework Program, No. 215056), 2008.

[27] Nick Mathewson and Roger Dingledine. 'Practical Traffic Analysis: Extending and Resisting Statistical Disclosure'. In: *Privacy Enhancing Technologies*. Ed. by David Martin and Andrei Serjantov. Vol. 3424. LNCS. Springer, 2005, pp. 17–34.

[28] Tim Moses. *eXtensible Access Control Markup Language (XACML) Version 2.0*. Standard. OASIS, 2005.

[29] Tim Moses. *Privacy policy profile of XACML v2.0*. Standard. OASIS, 2005.

[30] PETwebII partners. *Privacy-respecting identity management for e-Norge (PETwebII, funded by the Research Council of Norway in the VERDIKT program, no. 193030)*. `http://petweb2.projects.nislab.no/`. last accessed: April 04, 2011.

[31] PrimeLife partners. *PrimeLife: Bringing Sustainable Privacy and Identity Management to Future Networks and Services*. Project Deliverable D3.1.2. Privacy and Identity Management in Europe for Life (PrimeLife, Project within the European Community's 7th Framework Program, No. 216483), 2008.

[32] Andreas Pfitzmann and Marit Hansen. *Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management – A Consolidated Proposal for Terminology*. `http://dud.inf.tu-dresden.de/Anon_Terminology.shtml`. (Version 0.34). 2010.

[33] Andreas Pfitzmann and Marit Hansen. 'Anonymity, Unobservability and Pseudonymity – A Proposal for Terminology'. In: *Designing Privacy Enhancing Technologies, Proceedings of workshop on Privacy-Enhancing Technology (PET) 2000*. Vol. 2009. LNCS. Springer, 2001, pp. 1–9.

[34] György (George) Pólya. *How to Solve It: A New Aspect of Mathematical Method*. 2nd. Princeton University Press, 1957.

[35] Karl Raimund Popper. *The Logic of Scientific Discovery*. London: Hutchinson, 1959.

[36] Dave Raggett. *Draft 2nd Design for Policy Languages and Protocols*. Project Heartbeat 5.3.2. PrimeLife project of the European Community's 7th Framework Program, No. 216483, 2009.

[37] Joseph Reagle and Lorrie Faith Cranor. 'The platform for privacy preferences'. In: *Communications of the ACM* 42.2 (1999), pp. 48–55.

[38] Ron Rivest, Adi Shamir and Leonard Adleman. 'A Method for Obtaining Digital Signatures and Public-key Cryptosystems'. In: *Communications of the ACM* 21.2 (1978), pp. 120–126.

[39] Jan Schlörer. 'Zum Problem der Anonymität der Befragten bei statistischen Datenbanken mit Dialogauswertung [On the problem of respondents' anonymity in statistical databases with dialogue analysis]'. In: *4. GI-Jahrestagung*. Ed. by D. Siefkes. Vol. 26. LNCS. Springer, 1975, pp. 502–511.

[40] Andrei Serjantov and George Danezis. 'Towards an information theoretic metric for anonymity'. In: *Workshop on Privacy Enhancing Technologies*. Ed. by Paul Syverson and Roger Dingledine. Vol. 2482. LNCS. Springer, 2002, pp. 41–53.

[41] Claude Elwood Shannon. 'A Mathematical Theory of Communications'. In: *Bell System Technical Journal* 27 (1948), pp. 379–423, 623–656.

[42] Sandra Steinbrecher and Stefan Köpsell. 'Modelling Unlinkability'. In: *Workshop on Privacy Enhancing Technologies*. Ed. by Roger Dingledine. Vol. 2760. LNCS. Springer, 2003, pp. 32–47.

[43] Latanya Sweeney. '*k*-Anonymity: A Model for Protecting Privacy'. In: *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10.5 (2002), pp. 571–588.

[44] Hal Ronald Varian. 'Internet Policy and Economics – Challenges and Perspectives'. In: Published in Privacy and Self-Regualtion in the Information Age, a report issued by the NTIA, 1997. Springer, 2009. Chap. Economic Aspects of Personal Privacy, pp. 101–109.

[45] Alan Westin. *Privacy and Freedom*. New York: Antheneum, 1967.

Inter-temporal Privacy Metrics

Informational privacy of individuals has significantly gained importance after information technology has become widely deployed. Data, once digitalised, can be copied, distributed, and long-term stored at negligible costs. This has dramatic consequences for individuals that leave traces in the form of personal data whenever they interact with information technology, for instance, computers and phones; or even when information technology is recording the personal data of aware or unaware individuals. The right of individuals for informational privacy, in particular to control the flow and use of their personal data, is easily undermined by those controlling the information technology.

The objective of this thesis is to study the measurement of informational privacy with a particular focus on scenarios where an individual discloses personal data to a second party which uses this data for re-identifying the individual within a set of other individuals. We contribute with privacy metrics for several instances of this scenario in the publications included in this thesis, most notably one which adds a time dimension to the scenario for modelling the effects of the time passed between data disclosure and usage. The result is a new framework for inter-temporal privacy metrics.

ISBN 978-91-7063-603-5

ISSN 1403-8099

DISSERTATION | Karlstad University Studies | 2014:63