Privacy and Identity Management in Europe for Life

Second report on mechanisms

Editors:	Stefano Paraboschi (UNIBG)
Reviewers:	Dieter Sommer (IBM)
	Harald Zwingelberg (ULD)

Identifier:	D2.3.1
Type:	Deliverable
Version:	1.0
Class:	Public
Date:	February 28, 2010

Abstract

Today's society places great demand on the dissemination and sharing of information. One of the main challenges is to enable the legitimate use and sharing of information while at the same time guaranteeing both proper protection of the privacy of the individuals to whom information refers and proper preservation of the user's authority over the data.

This document describes the research results of the second year of the project obtained by the four work packages of Activity 2, which all focus on the investigation of the issues above. The document, similarly to the report produced at the end of the first year [CS09], includes one chapter for each work package that briefly describes the main research results along with an indication of what are the issues that will be addressed in the remaining third year of the project. The last chapter lists the abstracts of the research papers of the second year, reporting the findings of the work packages of Activity 2.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 216483 for the project PrimeLife.

Members of the PrimeLife Consortium

1.	IBM Research GmbH	IBM	Switzerland
2.	Unabhängiges Landeszentrum für Datenschutz	ULD	Germany
3.	Technische Universität Dresden	TUD	Germany
4.	Karlstads Universitet	KAU	Sweden
5.	Università degli Studi di Milano	UNIMI	Italy
6.	Johann Wolfgang Goethe - Universität Frankfurt am Main	GUF	Germany
7.	Stichting Katholieke Universiteit Brabant	TILT	Netherlands
8.	GEIE ERCIM	W3C	France
9.	Katholieke Universiteit Leuven	K.U.Leuven	Belgium
10.	Università degli Studi di Bergamo	UNIBG	Italy
11.	Giesecke & Devrient GmbH	GD	Germany
12.	Center for Usability Research & Engineering	CURE	Austria
13.	Europäisches Microsoft Innovations Center GmbH	EMIC	Germany
14.	SAP AG	SAP	Germany
15.	Brown University	UBR	USA

List of Contributors

Contributions from several PrimeLife partners are contained in this document. The following list presents the contributors for the chapters of this deliverable.

Chapter	**Author(s)**
Executive summary	**UNIBG**
Chapter 1: Cryptographic mechanisms (WP2.1)	**IBM**, GD, K.U.Leuven, UBR
Chapter 2: Mechanisms supporting users' privacy and trust (WP2.2)	**TUD**, KAU, TILT
Chapter 3: Privacy of data (WP2.3)	**UNIBG**, UNIMI, SAP, TILT
Chapter 4: Access control for the protection of user-generated data (WP2.4)	**UNIMI**, UNIBG, EMIC, SAP
Chapter 5: Conclusions	**UNIBG**
Chapter 6: Abstracts of research papers	**UNIBG**, IBM, UNIMI, TUD, SAP, UBR, KAU, K.U.Leuven, GD

Contents

Executive summary **9**

1 Cryptographic mechanisms (WP 2.1) **11**
1.1 Introduction . . . 11
1.2 Research results . . . 12
1.2.1 Cryptography for privacy and trust (Task 2.1.1) . . . 12
1.2.2 Trusted wallet (Task 2.1.2) . . . 22
1.3 Future Research . . . 23
1.3.1 Cryptography for privacy and trust (Task 2.1.1) . . . 23
1.3.2 Trusted wallet (Task 2.1.2) . . . 24

2 Mechanisms supporting users' privacy and trust (WP 2.2) **25**
2.1 Introduction . . . 25
2.2 Research results . . . 26
2.2.1 Transparency support tools (Task 2.2.1) . . . 26
2.2.2 Privacy measurement (Task 2.2.2) . . . 29
2.2.3 Privacy-respecting establishment of collaborative groups (Task 2.2.3) 32
2.2.4 Trust management by interoperable reputation systems (Task 2.2.4) 38
2.2.5 Privacy awareness (Task 2.2.5) . . . 41
2.3 Future research . . . 44
2.3.1 Transparency support tools (Task 2.2.1) . . . 44
2.3.2 Privacy measurement (Task 2.2.2) . . . 45
2.3.3 Privacy-respecting establishment of collaborative groups (Task 2.2.3) 45
2.3.4 Trust management by interoperable reputation systems (Task 2.2.4) 45
2.3.5 Privacy awareness (Task 2.2.5) . . . 45

3 Privacy of data (WP 2.3) **47**
3.1 Introduction . . . 47
3.2 Research results . . . 48
3.2.1 Privacy assessment and privacy metrics (Task 2.3.1) . . . 48
3.2.2 Techniques for enforcing data privacy (Task 2.3.2) . . . 49
3.2.3 Efficient organization and access to privacy-preserving data collections (Task 2.3.3) . . . 53
3.3 Future research . . . 54
3.3.1 Privacy assessment and privacy metrics (Task 2.3.1) . . . 54
3.3.2 Techniques for enforcing data privacy (Task 2.3.2) . . . 55

3.3.3 Efficient organization and access to privacy-preserving data collections (Task 2.3.3) 55

4 Access control for the protection of user-generated data (WP 2.4) 57
4.1 Introduction 57
4.2 Research results 58
4.2.1 Dissemination control and secondary use restrictions (Task 2.4.1) . 58
4.2.2 Access control to confidential data stored at external services (Task 2.4.2) 61
4.2.3 User-managed access control to personal data stored in trusted services (Task 2.4.3) 65
4.3 Future research 65
4.3.1 Dissemination control and secondary use restrictions (Task 2.4.1) . 65
4.3.2 Access control to confidential data stored at external services (Task 2.4.2) 65

5 Conclusions 67

6 Abstracts of research papers 69
6.1 Cryptographic mechanisms (WP 2.1) 69
6.2 Mechanisms supporting users' privacy and trust (WP 2.2) 74
6.3 Privacy of data (WP 2.3) 76
6.4 Access control for the protection of user-generated data (WP 2.4) 79

Bibliography 92

List of Figures

1 Components in the log system 28
2 The attacker's view on the system 30
3 Simulation results for 100 messages 32
4 Faster convergence with larger distances between clusters 33
5 Model of system environment 38
6 Binding between reputation granting and interactions 40
7 User interface of the forum 43

8 An example of relation (a) and of confidentiality constraints over it (b) . . 52

9 Enlarging (a), shifting (b), and reducing (c) 61
10 An example of access matrix (a) and of user graph over users $\{A,B,C,D\}$ (b) 63
11 A user tree (a) and the corresponding key rings (b) 64

List of Tables

2 2×2 design of the study 42
3 Perceived Privacy Index for different groups according to the presented privacy-awareness information 44

Executive summary

The deliverable presents the major results of the research done in Year 2 of PrimeLife in Activity 2. The goal of Activity 2 is to investigate the development of novel privacy mechanisms, considering a variety of scenarios and approaches.

Activity 2 is organized in four work packages that address the above research challenges. *Work Package 2.1: Cryptographic mechanisms* focuses on cryptographic techniques for supporting privacy and trust. *Work Package 2.2: Mechanisms supporting users' privacy and trust* focuses on solutions for supporting users in checking whether their personal data are used in accordance with privacy laws and privacy constraints specified by them. *Work Package 2.3: Privacy of data* focuses on solutions for assessing and ensuring privacy of large collections of sensitive data. *Work Package 2.4: Access control for the protection of user-generated data* focuses on solutions for enabling the enforcement of access restrictions on data generated from and/or by users.

The work produced results in the form of papers and software prototypes. Papers have been published on major international conferences and journals, thanks to the efforts by all the partners involved in Activity 2. Chapter 6 presents the abstracts of all the PrimeLife publications associated with the workpackages in Activity 2. Overall, Activity 2 was responsible in the second year of PrimeLife for the preparation of 32 publications, so classified: 3 papers published on international scientific journals, 25 papers published in the proceedings of international scientific conferences, 2 chapters in international scientific books, 2 chapters in the book prepared after the PrimeLife summer school.

Software prototypes have been developed by all the workpackages in Activity 2, which are going to contribute to the work in other activities (Activity 1, Activity 3, and Activity 5). Several of the tools will contribute to work package WP3.3 on open source, as they are planned to be released with an open source license and will be the subject of additional development in the third year of the project.

The main research results of each work package are summarized in the following.

- *WP2.1: Cryptographic mechanisms.* This work package has produced results in eight different areas: *1)* anonymous credential systems; *2)* private service access; *3)* delegation of access rights; *4)* group signatures; *5)* privacy-aware third party services; *6)* selective access in social networks; *7)* anonymous encryption; *8)* trusted wallets.

- *WP 2.2: Mechanisms supporting users' privacy and trust.* This work package has produced results in five different areas: *1)* mechanisms supporting transparency; *2)* privacy measurements for assessing the privacy of a user in communication systems; *3)* privacy-aware establishment of collaborative groups; *4)* trust management by interoperable reputation systems; *5)* mechanisms for supporting privacy awareness.

- *WP2.3: Privacy of data.* This work package has produced results in five different areas: *1)* design of support tools for anonymization; *2)* privacy metrics and privacy enforcing methods (anonymization) to assess the privacy protection exhibited by data collections; *3)* protection of statistical information; *4)* specification and enforcement of privacy constraints in relational databases; *5)* optimization of fragmented schemas.

- *WP2.4: Access control for the protection of user-generated data.* This work package has produced results in five different areas: *1)* XML secure views using semantic access control; *2)* context information and location privacy; *3)* access control to outsourced data; *4)* optimization of the representation of access control policies; *5)* management of personal information stored on trusted services.

The document is organized in four chapters, one for each work package (Chapters 1-4). Each chapter describes the research work performed in the second year of PrimeLife as well as work planned for the third year. Chapter 5 draws some conclusions on the work done in the second year of the project. Chapter 6 reports the abstracts of the research papers presenting the findings of the work packages throughout the reporting period.

Chapter 1

Cryptographic mechanisms (WP 2.1)

1.1 Introduction

Cryptographic schemes and protocols that support users in protecting their personal privacy gain in importance because most of our personal data is nowadays stored and processed in digital format. Encryption, hashing, and simple authentication techniques form the basis for secure information processing systems. However, more advanced cryptographic mechanisms are needed if we desire security properties that go beyond the simple confidentiality and authenticity properties. This holds particularly true, if no single party is trusted by all the users, and all users want to retain some control over their data. Such protocols include, but are not limited to, anonymous credentials [Cha85, CL01, Lys02], electronic cash [Cha82, Bra93, CHL05], group signatures [CvH91, CS97, ACJT00, BBS04], private information retrieval, oblivious transfer, blind signatures [Cha82], and, in the most generalized case, secure multi-party computation of any efficiently computable function.

This activity is concerned with research that distills and if necessary extends such mechanisms into a toolkit for solving privacy-enhancing identity management problems.

In today's electronic transactions, delegation of credentials is increasingly necessary. Without delegation, only the most stringent access control policies can be realized. Those that are actually in use today do require delegation. Delegation captures trust relationships on which communities are built. Lack of efficient cryptographic algorithms for delegation of anonymous credentials is an impediment for adopting anonymous credentials in practice, and stands in the way of making privacy-preserving protocols the default.

Furthermore, the scenarios of PrimeLife's Activity 1 raise a number of open cryptographic problems for which no solutions exist. PrimeLife is concerned with issues such as social networks privacy, long-lived privacy, and the deployment of privacy mechanisms in the real world through the means of open source initiatives and standardization. Mechanisms such as delegation of credentials, searchable encryption, and oblivious service

access, prove useful within the thick social fabric of today's Internet where private information need not only be protected but, in addition, needs to be processed and shared in a responsible manner.

Examples for such mechanisms are signature scheme that allow one to copy only parts of a signed text, such that the recipient can still be convinced that the part originated from an anonymous but trustworthy source. Another example are special encryption and key management schemes that allow users to enforce, in a distributed fashion (i.e., without any trusted components such as a central server or DRM-like hardware), that their data can only be read by specific users, e.g., their friends.

Work of this work package is divided into two tasks. The first one, *Task 2.1.1: Cryptography for Privacy and Trust*, is concerned with the study and design of cryptographic algorithms for privacy and trust. The second task, *Task 2.1.2: Trusted Wallet* is focused on the design and development of components that allow users to securely manage their cryptographic key materials such as secret keys and electronic credentials.

1.2 Research results

1.2.1 Cryptography for privacy and trust (Task 2.1.1)

In the reporting period, work package WP 2.1 has improved the state of the art in the following major areas.

Anonymous Credentials We have continued our research on new credential systems to improve both the efficiency of the attribute encoding and the revocation of credentials. We have also investigated the standard-model security of Schnorr signatures, which are at the basis of the U-Prove anonymous credential system.

Delegation While the delegation of non-anonymous credentials is easily solved and widely used in the form of a certification hierarchy, a solution to this problem for anonymous credentials proved to be evasive. We published a solution to this long standing problem in [BCC+09] and continue to look at accountability mechanisms for delegatable anonymous credentials.

Group Signatures We started investigating a new design paradigm for group signature schemes. This resulted in very short and efficiently computable signatures.

Private Service Access In certain sensitive cases users may be interested in using a service unobservably, i.e., even the service provider would not learn the exact type of service requested. This is important in order to protect against insiders, especially if the service itself is dependent on the personal information of other data subjects, as it is the case for social networking sites. We investigated ways to incorporate privacy-friendly service selection and payment protocols into such a system. In particular, we provide protocols for searching on encrypted data and for priced oblivious transfer. From a more theoretical angle, we are looking for oblivious transfer protocols that fulfill the strongest notions of security for the composition of multiple protocols.

We also looked at two extensions for priced oblivious transfer. The first allows for anonymous and unlinkable purchases, while the second aims at protecting both the buyer and the merchant against unfair behavior of the other party. For example, a merchant could try to refuse finishing the transfer of a digital good and nevertheless keep the payment.

Privacy Aware Services Offering services in a privacy friendly way is often very challenging and requires new cryptographic mechanisms and protocols. The goal here is to achieve privacy while maintaining accountability, by using new protocols but also by possibly relying on third parties to enforce security properties. Examples of such parties could be identity escrow services. Thereby, these third parties should not be involved in the normal operation of the system, but only in case some party misbehaves. Third parties play a role in securely complying to legal requirements, such as data retention.

The trust that other parties need to put into such third party should be minimal. Ideally, the third party should be completely accountable, i.e., everyone should be able to verify that the party behaved as specified without being able to derive any other information. The work package studies different ways to offer services and how third parties can best be employed. First results here have been published.

Cryptography for Selective Access Control in Social Networks The aim of cryptographic support for selective access control in social networks is to give the user control not only about the definition of the access control policy, but also its enforcement. In a first step, we will develop policy concepts that are general enough to describe access control restrictions for a variety of different social networking sites. We will also investigate different means for enforcing such policies using sticky policies and advanced encryption techniques. The establishment of cryptographic trust between users in the form of adequately distributed keys is another big challenge in this area that we are working on.

Anonymous encryption Anonymous systems encryption does not only hide the content of the encrypted message, but also the intended recipient of the message. This can for instance improve the privacy of social networking users. One problem, however, is that the intended recipient himself should be able to detect messages directed to him. We investigate the "robustness" property that the underlying encryption scheme has to satisfy in order to guarantee correct and secure delivery of messages. Interestingly, it turns out that folklore solutions, such as encrypting redundancy, do not work in general. We provide generic transforms for public-key and identity-based encryption that confer the robustness property for schemes that do not have it naturally.

We also investigate protocols for anonymous identity-based encryption schemes that have applications to private searching on encrypted data.

The results of our work are described under separate subsections below.

Anonymous credentials

A *credential system* is a system in which users can obtain credentials from organizations and demonstrate possession of these credentials. For instance, such a credential could be a driver's license containing as attributes data about the user such as her birthday, address, or the date she took the driving test. Other examples include passports, identity cards, or educational certificates. That is, a credential is a signed statement issued by some party (issuer) about another party (recipient or user).

Unlike a traditional PKI, in an *anonymous credential system*, such statements can be presented to third parties in a way such that the issuing and the presenting transaction are not linkable. Moreover, if the user wants to use her driver's license to show that she is of age, a credential system allows her to do so without revealing any other information about her. In other words, such systems allow the user to enforce what information another party learns about her. Thus, such credential systems form the foundation of any privacy enhancing identity management system and hence one of the goals of the work package is to improve the state of the art in this space. This includes making the existing protocols more efficient, realizing new features as to enable new application, and providing schemes that are based on simpler cryptographic assumptions.

Progress in the Second Project Year. In the second project year, we have achieved a breakthrough result in the area of anonymous credential systems: for the first time we were able to show that anonymous credentials can be practically implemented on a standard Java smart card. On a more theoretical level, we investigated the standard-model security of the signature scheme underlying the U-Prove anonymous credential system.

Java Card Implementation Secure identity tokens such as electronic identity (eID) cards are emerging everywhere. However, when implementing an anonymous credential system on inexpensive hardware platforms, which are typically used for eID cards, these schemes could not be made to meet necessary requirements, such as future-proof key lengths and transaction times on the order of 10 seconds. The main reason for this is the need for the hardware platform to be standardized and certified. Therefore an implementation is only possible as a Java Card applet. This results in severe restrictions: little memory (transient and persistent), an 8-bit CPU, and access to hardware acceleration for cryptographic operations only by defined interfaces, such as RSA encryption operations.

Still, we managed to realize the first practical implementation of an anonymous credential system on a Java Card 2.2.1. We achieved transaction times that are orders of magnitudes faster than those of any prior attempt, while raising the bar in terms of key length and trust model. Our system is the first one to act completely autonomously on card and to maintain its properties in the face of an untrusted terminal. In addition, we provided a formal system specification.

Security of Schnorr signatures We provided two necessary conditions on hash functions for the Schnorr signature scheme to be secure, assuming compact group representations such as those that occur in elliptic curve groups. We also show, via an

argument in the generic group model, that these conditions are sufficient. Our hash function security requirements are variants of the standard notions of preimage and second preimage resistance, mandating a use of n-bit hash functions to achieve a 2^n security level. Moreover, our analysis does not reveal any significant difference in hardness between forging signatures and computing discrete logarithms, playing down the importance of the loose reductions in existing random-oracle proofs, thereby supporting the use of "normal- size" groups.

Private service access

We considered protocols that allow users to access services without revealing their access patterns. Such techniques are known as oblivious transfer (OT) and Private Information Retrieval (PIR).

Oblivious access to protected resources could for instance be used to realize social networking services in which the underlying social network is itself hidden from the service provider, while users can still find out information about the friends of their friends. Another scenario is sensitive databases, e.g., DNA databases. A user that contributed his DNA to a DNA database has a right to be notified about every use of his data. However, the fact that a medical institute accessed a certain piece of DNA may already be extremely sensitive and may need to be hidden even from the administrator of the database. The issues get more contrived as service providers need to control access to the database, e.g., to protect the privacy of the data subject or to selectively charge money for the access to certain data.

Progress in the Second Project Year. We investigated the case of "priced" oblivious transfer (POT), where a different price is associated with each message in the OT. The receiver and sender interact in such a way that the receiver pays the appropriate price for the selected message, but the sender does not learn which message was selected, nor the price that was paid. Both OT and POT admit an adaptive variant where the receiver can choose a new message index σ_i after having received $m_{\sigma_{i-1}}$. Adaptive POT is the most promising variant for real-world applications such as privacy-preserving e-commerce.

We designed an adaptive OT scheme and we further modified it to construct an adaptive POT scheme. Both constructions are universally composable, optimal in terms of rounds of communication and, after an initialization phase of complexity $O(N)$, both have constant communication and computational cost in each transfer phase. This work was published this year [BCGS09]. In follow up work, we are looking at how to guarantee the fairness of priced oblivious transfer with the help of techniques from optimistic fair exchange. Here a third party is invoked as a judge to mediate between the receiver and the sender to guarantee that the receiver pays for what he receives and that the sender sends what he is payed for. The third party does not need to be always online and only is needed in case of conflict.

An adaptive POT protocol has to maintain state information to keep track of how much money each user has left on her account. In all protocols prior to our work, the sender kept (encrypted) state information for each user. This makes all transfers from the same user linkable, hence partially defeating the privacy goal of OT. We designed an adaptive POT scheme where it is the users themselves who maintain state informa-

tion, so that different transfers by the same user remain unlinkable from the sender's point of view. Moreover, our protocol offers the possibility to recharge the amount of money in a user's account while hiding the account balance from the sender. Previous protocols did not offer such functionality, so that users had to choose between losing the remaining money in their accounts, or revealing it to the sender, thereby possibly leaking information about their purchased messages.

Although privacy properties are guaranteed, current schemes do not offer fair exchange. A malicious vendor can, e.g., prevent the buyer from retrieving the goods after receiving the payment, and a malicious buyer can also accuse an honest vendor of misbehavior without the vendor being able to prove this untrue.

In order to address these problems, we defined the concept of optimistic fair priced oblivious transfer and propose a generic construction that extends secure POT schemes to realize this functionality. Our construction, based on verifiably encrypted signatures, employs a neutral adjudicator that is only involved in case of dispute, and shows that disputes can be resolved without the buyer losing her privacy, i.e., the buyer does not need to disclose which digital goods she is interested in. We showed that our construction can be instantiated with an existing universally composable POT scheme, and furthermore we proposed a novel full-simulation secure POT scheme that is much more efficient.

Delegation

A recurring problem in identity management systems is the delegation of access rights. Delegation is a common tool in the physical world, when one lends to someone else one's keys or when one signs a paper document that delegates certain rights (for example to vote, to buy a house, or to submit a proposal). In the on-line world this can be achieved by passing on a password, but it is clear that this brings unintended consequences, like the difficulty in limiting its use. Moreover, there are many complex ways of delegating access rights or credentials: there are delegations with and without restrictions in scope, with and without anonymity, delegations that can be passed on themselves etc.

Progress in the Second Project Year. Conventional public key certificates support delegation. A certification authority can certify a company, and the company can in turn certify its employees. These certificates do however leak the identity of all intermediary certificates.

Traditional anonymous credentials do not support delegation, and the implementation of delegation properties similar to those described above are difficult to achieve. We provide some first results to allow for anonymous credentials that allow to prove statements such as the following: "My company has a subscription to your database, and has granted me access to a specific subset of the database", without revealing the name of my company.

In [BCC+09], published in the second project year, we revised the entire approach to constructing anonymous credentials and identify *randomizable* zero-knowledge proof of knowledge systems as the key building block. We formally define the notion of randomizable non-interactive zero-knowledge proofs, and give the first construction by showing how to appropriately rerandomize Groth and Sahai (Eurocrypt 2008) proofs. Our insight is that instead of giving Alice his signature, Oliver gives Alice a non-interactive proof-

of-knowledge of the signature. The trick is to find a proof-system that would then let Alice (1) delegate the credential by extending the proof and (2) rerandomize the proof every time she shows (or extends it) to preserve her anonymity.

An additional requirement in many applications of delegation is that if Alice delegates a right to Bob, she at the same time loses the right she is delegating. This is related to transferable e-cash, in which a user that transfers a coin to someone else cannot spend the coin anymore. We are looking at solutions for transferable e-cash that can also be used for revoking the rights of delegating users. Previous constructions achieving strongly anonymous transferable e-cash were merely of theoretical interest because of their inefficiency. In addition, we want to achieve two-sided anonymity, i.e., the receiver is also anonymous to the spender.

Group signatures

Group signatures were introduced in 1991 by Chaum and van Heyst [Cv91]. Such schemes allow members of a group to anonymously sign a message on behalf of the whole group. As an example application, they allow an employee of a company to sign a document in such a way that the verifier gets only to know that it was signed by an employee, but not the particular employee involved. Control of the group membership is provided by a *Group Manager*, who can add users (called *Group Members*) to the group. In addition, there is an entity called *Opener* who can reveal the identity of signers, which is necessary in the case of disputes. In some schemes, the two tasks of adding members and revoking anonymity are combined into a single role.

Since 1991 a number of security properties have been developed for group signatures including unforgeability, anonymity, traceability, unlinkability, and non-frameability. In 2003 Bellare, Micciancio and Warinschi [BMW03] developed what is now considered the standard security model for group signatures. They propose two simple security properties called *full anonymity* and *full traceability* and show that these capture the previous security requirements of unforgeability, anonymity, traceability, and unlinkability. One property not covered by their notion is *non-frameability* (or exculpability), in which the Group Manager cannot produce a signature that can be falsely attributed to an honest Group Member.

Schemes based on bilinear maps are the most efficient known group signatures, both in terms of bandwidth and in terms of computational efficiency. Of particular importance being the schemes of Boneh, Boyen and Shacham [BBS04] (BBS), Boneh and Shacham [BS04] (BS), Camenisch and Lysyanskaya [CL04] (CL), Delerablée and Pointcheval [DP06] (DP), and Shacham [Sha07] (S).

Progress in the Second Project Year. We introduced a group signature scheme based on Type-3 pairings as described in [GPS08]. That is, we have cyclic groups $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ of prime order. There exists a bilinear map $\hat{h} : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ with $\forall x \in \mathbb{G}_1$, $\tilde{y} \in \mathbb{G}_2$ and $\alpha, \beta \in \mathbb{Z}_q$ we have $\hat{h}(x^\alpha, \tilde{y}^\beta) = \hat{h}(x, \tilde{y})^{\alpha\beta}$ and $\hat{h}(g, \tilde{g}) \neq 1$.

We make use of the LRSW assumption as introduced by Lysyanskaya et al. [LRSW99] to prove the non-traceability of our scheme. In addition we use a bilinear version of the standard DDH assumption, namely the XDDH assumption, which holds if DDH is hard in the pairing group $\mathbb{G}_1$. To demonstrate the non-frameability of our scheme we require

an additional assumption, which we call the symmetric Discrete Logarithm Assumption. The latter states that given the tuple $(g^\mu, \tilde{g}^\mu) \in \mathbb{G}_1 \times \mathbb{G}_2$ computing μ is a hard problem.

Intuitively, our scheme is based on two special properties of CL signatures, namely on their re-randomizability and on the fact that the signature "does not leak" the message that it authenticates. Intuitively, a user's group signing key is a CL signature on a random message ξ that only the user knows. To create a group signature for a message m, the user re-randomizes the CL signature and attaches a signature proof of knowledge of ξ on m. In a simplified version of our scheme, group signatures could be opened by letting the Group Manager check for which of the issued values of ξ the re-randomized CL signature is valid.

To obtain non-frameability, however, in our scheme the Group Manager does not know ξ itself, but rather ξ is generated jointly during an interactive Join protocol between the user and the Group Manager. Essentially, this protocol is a two-party computation where the user and the Group Manager jointly generate ξ, a valid CL signature on ξ, and a key derived from ξ that allows the Group Manager to trace signatures, but not to create them.

We prove the anonymity, traceability and non-frameability of our scheme. *Anonymity* requires that group signatures do not reveal the identity of the signer. In contrast to the standard notion as described in [BMW03], we allow a user to trace his own signatures. We say that the group signature scheme satisfies the anonymity property if for any probabilistic polynomial-time adversary, the advantage in identifying a user's signature is a negligible function of the security parameter.

Informally, *traceability* requires that no adversary can create a valid signature that cannot be traced to some user that had already been registered. We model the strong but realistic setting where all of the signers are corrupt and work against the group manager. We say that the scheme is traceable if for any probabilistic polynomial-time adversary, its advantage in creating a non-traceable signature is a negligible function of the security parameter. Our notion of traceability differs from the traditional one in that, according to our definition, an attacker that creates a signature that opens as some honest identity is not considered an attack.

Non-frameability requires that even a cheating Group Manager cannot falsely accuse an honest user of having created a given signature. We say that a scheme is non-frameable if for any probabilistic polynomial-time adversary the advantage in framing a user is a negligible function of security parameter.

Privacy aware third-party services

Services with identity management relations have a challenging set of requirements. They require accountability of identity data as well as privacy protection of their users. The first requirement is usually driven by the need to mitigate risks to the service provider, and the latter requirement is driven by regulatory compliance (e.g., privacy legislation [Eur95]). This requires new primitives for accountable and privacy-preserving identity management methods for services that provide improved support for these two requirements.

Anonymous credential systems exist that provide methods to achieve a wide range of accountability and privacy goals for services. In particular anonymous credentials

systems already have the capability to selectively disclose statements about a user, e.g., proving the user's age without revealing the user's actual date of birth. With the complementary primitive of verifiable encryption [CS03], credential systems can also provide accountability without infringing on privacy requirements. For instance, a user U can encrypt her true identity to the authorities, provide this encrypted data to a service provider SP, and convince SP in a zero-knowledge proof of knowledge that this encrypted data contains a valid user identity that can be opened by the authorities.

Existing systems realize accountability with verifiable encryption by encrypting the user's identity to a trusted third party (TTP), for instance, law enforcement authorities. This presents three challenges:

1. This mode of operation involves a fully trusted TTP with no graceful degradation of privacy and security should the TTP become compromised.

2. A malicious service provider holding a verifiable encryption may attempt to betray the user by opening a decryption case at the TTP without good cause.

3. Honest service providers find the traditional system encumbering because of the need to involve such highly trusted authorities for even minor dispute cases. For example, to bring a case to law enforcement in the real world is likely to have a non-trivial cost, both in the time required, and in support from legal council.

Therefore, we propose to rethink the corresponding primitives to better meet the actual needs of services.

Progress in the Second Project Year. Legislation on data retention, including EU data retention direction (Directive 2006/24/EC), requires that communication service providers such as internet service providers (ISPs) and telecommunication providers (telcos) store user-related data over an extended period of time (from 6 to 24 months), so that it is available for criminal investigations and law enforcement. While these laws also require such data to be adequately protected, practice shows that security breaches are common when data is stored at such a large scale.

Although storing such data in encrypted form is an obvious protection method, this is not always an effective solution. The data must also be made available so that it can be searched, accessed, and analyzed. Decrypting the entire database for each search is not practical and decreases security. In particular, it provides no protection against malicious insiders. The party doing the search may not (or should not) have the authority to see the whole database.

We describe a solution [CKRS09] for an application scenario in which a communication service provider retains data in compliance with the EU data retention directive. Our solution achieves the following properties: (1) The data is stored in a secure and efficient manner. (2) An investigator is able to retrieve data relevant to the specific (authorized) investigation while still protecting irrelevant (unauthorized) data. (3) Neither the data holder nor the key holder learn any details about the investigation. The ISP's routers and database farms are data holders, while an independent third party or the ISP's high security data center is the key holder.

Cryptography for selective access control in social networks

Anonymous credentials are not a perfect fit for the privacy needs of social network (SN) sites. The users of social networks do not aim at data minimization. They are using it to share information. As we will see, there will still be some legitimate uses of anonymous credentials in social networks, but a different core strategy is needed.

The main privacy concern of SN users is not data minimization but control. Users want to specify who can see which information (specify access control policies) and have them enforced by the SN, preferably without putting too much trust in external parties (everyone who is not them, or their peers).

The main mechanism to enforce access control is a reference monitor. The main mechanism to reduce trust in external parties is to (1) either store data locally (or at a proxy that they trust), which would require users to be always online (or invest in additional infrastructure), or (2) store data in encrypted form. We consider the latter approach. The reference monitor and the encryption need to be integrated with each other in an adequate way.

Progress in the Second Project Year. We are currently collaborating with Activity 1 to extract additional requirements for advanced cryptographic techniques for policy enforcement. However, one first needs to exploit the full power of classical hybrid encryption techniques, before it becomes reasonable to start new research in this area.

An orthogonal subject is trust establishment and trust management among users. The establishment of a PKI where each entity has its own public key pair could lead to a substantial increase in security in today's Web infrastructure. As the approach of letting users cross-certify their certificates, known as the Web-of-trust, did not lead to a wide spread adoption, we investigated an alternative solution.

In [BMP+09] we propose to use relationships established through one or several electronic social networks (ESN) as a means to automatically suggest cross-certification of certificates. As an example, an ESN is aware of two users communicating often and regularly. Given that both users uploaded a (traditional) certificate with the same personal information as they use within the ESN, the latter might suggest that the users mutually cross-certify their certificates. We propose a number of mechanisms that allow an intuitive trust establishment process. In addition, the overhead of the trust management is kept at a minimum as it is closely aligned to the management of the ESN.

We are currently investigating how to combine three approaches to get the best feature mix for the Activity 1 prototype: establishment of individual trust through an out-of-band channel between two parties (this is the most intuitive, but also most cumbersome way for establishing trust), traditional cross certification, i.e., a web of trust, and various ways for improving cross certification using an ESN.

Anonymous encryption

The primary security requirement for encryption is data privacy, as captured by the security notions of indistinguishability under chosen-plaintext (IND-CPA) and chosen-ciphertext (IND-CCA) attacks. Increasingly, we are also seeing a market for *anonymity*,

asking that a ciphertext does not reveal the encryption key under which it was created.

We investigated a provable-security treatment of a third security notion of encryption, called "robustness." Robustness reflects the difficulty of producing a ciphertext valid under two different encryption keys. Robustness helps make encryption more misuse resistant.

Where you need anonymity, there is a good chance you need robustness too. Indeed, we would go so far as to say that robustness is an essential companion of anonymous encryption. The reason is that without it we would have security without basic communication correctness, likely upsetting our application. This is best illustrated by the following canonical application of anonymous encryption, but shows up also, in less direct but no less important ways, in other applications. A sender wants to send a message to a *particular* target recipient, but, to hide the identity of this target recipient, anonymously encrypts it under her key and broadcasts the ciphertext to a larger group. But as a member of this group I need, upon receiving a ciphertext, to know whether or not I am the target recipient. (The latter typically needs to act on the message.) Of course I can't tell whether the ciphertext is for me just by looking at it since the encryption is anonymous, but decryption should divulge this information. It does, unambiguously, if the encryption is robust (the ciphertext is for me iff my decryption succeeds) but otherwise I might accept a ciphertext (and some resulting message) of which I am not the target, creating miscommunication.

Anonymous encryption finds important applications in searchable encryption, where encrypted messages can be searched for keywords without decrypting the message. In a clever usage of anonymity, Boneh et al. [BCOP04] showed how this property can turn an identity-based encryption scheme into a searchable encryption scheme. But Abdalla et. al [ABC$^+$08] noted that their construction could lack *consistency*, meaning turn up false positives. Abdalla et al.'s solution was to modify the construction, but what we observe instead is that consistency would in fact be provided by the *original* construct if the IBE scheme was robust.

Progress in the Second Project Year *Robustness* is trivially added to any encryption scheme by concatenating the intended recipient's identity or public key to the ciphertext. This solution is not an option for anonymous encryption however. We found that natural "solutions" to providing robustness without violating anonymity, such as including the encryption key or identity of the target recipient in the plaintext before encryption and checking it upon decryption, do not work.

Apart from considering and dismissing the natural approaches that do not work, we provided formal definitions of several variants of robustness, and provided two general robustness-adding transforms. Moreover, we tested the robustness of existing schemes, including DHIES, the Cramer-Shoup scheme, and the generic Fujisaki-Okamoto and Canetti-Halevi-Katz transforms. It turns out that some of these are inherently robust, some can be made robust with minor modifications, and for others one has to apply our generic transformations to confer robustness.

In [CKRS09] we also designed a *committed blind anonymous identity-based encryption (IBE)* scheme based on the anonymous IBE scheme due to [BW06]. As the scheme in [BW06] is only selective ID secure, we extended it with adapted ID security and proved the modified scheme secure. For the modified scheme, we designed a blind key

extraction protocol. This leads to the first blind anonymous IBE scheme we are aware of. We extended the definition of blind IBE to allow for the derivation of a secret key for a committed identity. We used commitments to enforce restrictions on the blinded keywords. This can ensure that a key is issued only to an individual who can prove (in zero-knowledge) that the identity used to compute the commitment fulfills certain conditions imposed by the key generation server.

1.2.2 Trusted wallet (Task 2.1.2)

An identity management wallet has a similar function as a wallet in the real world. Instead of authenticating towards the different service providers, the user would authenticate to her "trusted identity management wallet", which contains all of her access credentials. By concentrating the sensitive data within a single application the potential harm is increased. We identified three compartments of an identity management system that should be protected independently:

1. The compartment running the applications that eventually want to make use of the IdM system.
2. The compartment doing the processing of the personal information (i.e., the wallet) and the interactions with the user.
3. The compartment containing the high security cryptographic material.

It is preferable to run the different compartments of a trusted wallet separated from each other and with different requirements on their security:

Low security: Application software such as the user's browser should run with normal operation system security. This guarantees good support of legacy software and good integration into the user's familiar computing environment.

Medium security: The wallet software including the identity selector is run in a compartment provided by a separate virtual machine.

High security: A smart card or a threshold scheme that remains secure as long as an adversary does not control a majority of a user's portable devices can be used to secure important cryptographic material.

Progress in the Second Project Year. The role of the high level security compartment is to act as a last trust anchor for the user's identity management system. If the user's computer gets hacked, if he loses his laptop, wrote all his passwords on a sheet of paper, or shared his whole wallet with an untrusted person or all his peers in a global file sharing network, the high security compartment will try its best to still provide some protection of a user's credentials and personal data.

This protection will not be perfect, a user can always decide to publish any information about himself over unprotected channels, and he can decide to act as an online proxy that shows credentials and allows everyone to impersonate him.

However, the high security compartment can still do its best to protect users who do not decide to shoot themselves into the foot. This indicates a second challenge, a

trade-off between security and usability. Adding security requires additional user interaction, e.g., to ask the user whether he agrees to a specific treatment of data or to enter a PIN. If the burden for the user becomes too high, he is likely to stop using the system. We did investigate interfaces that allow for an intuitive, lightweight representation of anonymous credentials. The visualization of individual attribute disclosure, that is, (zero-knowledge) proofs about attributes poses issues. We showed how established concepts of user interface design can be employed to help users familiarize themselves with these formerly unfamiliar identity management concepts. A remaining challenge is the integration of the additional security into the current design.

For the implementation of the high security compartment we experimented with porting anonymous credentials to a Java Card as described in Section 1.2.1. We managed to realize the first practical implementation of an anonymous credential system on a Java Card with transaction times that are orders of magnitudes faster than any prior attempt. However, to use our solution in the architecture of the trusted wallet requires further security provisions. This follows from the fact that the Java Card does not provide input/output capabilities.

In addition to implementing a full anonymous credential scheme on a Java Card, we experimented with implementing only the trusted platform model part of a direct anonymous attestation protocol on a Java Card. We also tested this functionality in a demonstrator that implements an electronic petition system, such as the one described in [DKD+09].

1.3 Future Research

1.3.1 Cryptography for privacy and trust (Task 2.1.1)

Anonymous Credentials. There are a large number of issues to solve to make anonymous credentials more versatile and applicable to a wider range of applications. Issues include the following:

Alternative variants for efficient revocation. Depending on the actual scenarios, the requirements for revocation differ much more than one would expect. We will continue to study different scenarios and their requirements and will try to come up with revocation mechanisms that are targeted at these scenarios.

Lightweight devices. We are going to study how to best achieve the functionality of private credentials for lightweight or restricted devices such as smart cards, as using the standard anonymous credential protocols (at least as soon as more than just the very basic functionality is needed) seems just a bit off from what would be accepted by end users in terms of computation times.

Private Service Access. The work in this area has just started and we will continue our work as outlined earlier. One interesting extension to the current approach would be to support stored procedures. Instead of retrieving a single data item, the service would compute a function of several data items, without learning which function was evaluated on which inputs. This corresponds to the private evaluation of a stored procedure. An example for such an application would be to compute similarities between genetic

material in order to establish ancestral relationships or the existence of a hereditary disease.

The recent discovery of fully homomorphic encryption [Gen09] is an important theoretical result that promises new results in this area.

Privacy Aware Third-Party Services. Also in this area we have so far only achieved initial results and hence plan to further investigate different scenarios and come up with new or better solutions. The end goal here is to derive a set of standard mechanisms that will allow one to implement any such service in a way that the third party is, on one hand, accountable, i.e., all other parties can verify that it has performed its task as expected and, on the other hand, is not aware of the identities of the parties involved in the transactions. The latter will ensure that the third party will not be able to change its behavior depending on the specific user involved.

For the data retention scenario, it would be useful to separate the responsibilities of third parties, e.g., one party could be responsible for enforcing the expiration of retention while another party verifies the validity of search requests signed by judges.

1.3.2 Trusted wallet (Task 2.1.2)

We will investigate further in the high security compartment. This might include hardware, such as a tamper resistant device with I/O capabilities, or new methods that extend the guarantees for the user. In addition, we will further investigate on how to visualize the security improvements to attain a usable solution. The prototype will show how the most critical requirements can be met along with how a trusted wallet can be employed to accomplish a smooth user experience.

Chapter 2

Mechanisms supporting users' privacy and trust (WP 2.2)

2.1 Introduction

The Internet offers its users numerous possibilities to interact with each other. Interactions cover various fields of interest and parts of life for many people. Examples most of us are familiar with are e-shopping, e-health, on line community services and e-government. Protecting *privacy* is an important issue in electronic interactions.

Here *transparency support tools* play an important role. They can be used either in addition or as a part of a privacy-enhancing identity management system (PE-IMS). Transparency is a legal privacy principle, which also can be derived from the EU Data Protection Directive 95/46/EC [Eur95]. When a data controller is requesting personal data from a data subject, the data controller must inform the data subject about the controller's identity, the purposes of data processing, the recipients of the data and all other information required to ensure the processing is fair ([Eur95] Art. 10). The data subject has the right to access all data processed about her, to demand the rectification, deletion or blocking of data that is incorrect or is not being processed in compliance with the data protection rules ([Eur95] Art. 12). The user's right to access also includes the right to obtain knowledge of the logic involved in any automatic processing of data concerning her. Even though there is no legal requirement that users can exercise their rights on the Internet, we believe that such a state of affairs would be beneficial for all parties involved and could also make the process administratively more efficient. Our efforts this year have been on constructing a privacy-preserving secure log system that with a proper support on the client side can be used by a data subject to securely and privacy-friendly access information on how her data has been used and processed. This log system has been implemented as a stand-alone tool [Ste09c] and has also been integrated in the PRIME system. This work is further described in Section 2.2.1

If one defines privacy as actively experienced "right to select what personal information about me is known to what people" [Wes67], privacy awareness encompasses a

users' perception, cognition and attention on whether others receive or have received personal information about her, her presence and activities, which personal information others receive or have received, who receives or has received personal information, and how these pieces of information are or might be processed and used. After working on the theoretical background [Ste09b, CS09] of *privacy awareness* in the first project year, in the second year we focused on developing a first version of a tool to support privacy-awareness and empirical research in order to study the effect of privacy-awareness information on users. While the tool was already presented in [Ste09c], the results of the empirical study are presented in Section 2.2.5.

Transparency and awareness need *privacy measurement*, especially a user's level of anonymity and the unlinkability when using services. There have been made several proposals to formalize and measure anonymity and unlinkability for both communication systems and applications. In Section 2.2.2 we elaborate on how much privacy of users is affected when combining linkability information between messages with information about senders of messages gained from the network layer.

Interaction systems usually both create and implement a virtual community [Rhe93]. Sharing personal information with others is the basic idea of many social applications like communities to create *trust* among users. Privacy enhancements tend to hinder trust, because trust usually is built up on information about and from the interaction partners distributed.

One known means for trust management is the use of reputation systems that manage information about past behavior of interaction partners. Based on this information, individuals can get a clue on how others might interact in the future. While in the first project year we focused on using convertible credentials [Ste09b, CS09] for building *privacy-respecting interoperable reputation systems* in the second project year we implemented a first version of a tool using this approach, as documented in [Ste09c]. In our theoretical work we concentrated on using anonymous payment systems as an alternative approach, as outlined in Section 2.2.4.

When introducing strong privacy mechanisms into Internet communities, interaction or even establishment of collaborative groups might become more difficult. One example is role management that cannot be predetermined, but a policy-credential-based mechanism is needed for the purpose of providing authentication and authorization. Another example is event scheduling without revealing the availability patterns of individuals to others users. In Section 2.2.3 we outline how role management and event scheduling can be done for a *privacy-respecting establishment of collaborative groups*.

2.2 Research results

2.2.1 Transparency support tools (Task 2.2.1)

Transparency enables data subjects to view what data is stored about them and how that data has been processed and used. In order to track usage and processing of data some form of audit trail is needed. This trail is commonly established in computer systems by using logs and logging systems. In order to be useful as a reliable source of information, this logging system would be expected to be protected against unauthorized modification and access. As the logging system in this case is used as a transparency tool for data

use and processing, it is itself a collection of processed personal data. Hence, we have to prevent that personal data is leaked to unauthorized users. In effect this means that the log needs to be a privacy preserving secure logging system. A number of secure logging systems exists in the literature (e.g., [SK98, MT09, Hol06, SSA06]), many of them based on and extending the Kelsey-Schneier log [SK98]. This type of log takes care of the integrity and confidentiality aspects of the log but does not solve all of the privacy preserving aspects of it. In particular, it does not provide unlinkability of log entries and secure anonymous access to log entries. Some work of unlinkability in connection with logs have been addressed by [WSLP08]. However, this work primarily addresses the unlinkability of logs between logging systems in an eGovernment setting rather than unlinkability of log entries within a log. Further, they do not address the problem of an inside attacker or provide anonymous access to log entries. Because of this we decided to design and implement a privacy preserving secure log for transparency purposes and incorporate it in the PRIME core. For an overview of the PRIME system see [CCKS07]. The log in itself is an extension of the Kalsey-Schneier log adding primarily unlinkability properties and the possibility of anonymous access to it. Detailed descriptions of the design of the log system is presented in [HPHL09, Ste09c]. This section will only give a brief description of the log and the implementation process.

Requirements

The requirements the log has to fulfill include the protection against unauthorized access to the log data. Since the log is primarily to be used as a transparency tool, the only entity authorized to access the clear text log data, in the most strict case, is the data subject herself or possibly her proxy. Based on this, requirements of the log will be the following:

- It should not be possible for anybody except the data subject to decrypt log entries once they are committed to the log.
- It should neither be possible to alter nor remove without detection entries made prior to an attacker taking control of the data controller's system.
- It should not be possible to link more than one log entry in the log referring to a specific data subject with that data subject except by the data subject herself. Ideally we would like to make it impossible to link any entry. However, our current solution makes one entry per data subject identifier linkable to that data subject identifier.
- For efficiency reasons the solution should as far as possible not require that the whole log database is fully traversed by any entity or sent as a whole to the data subject.

Structure of the Logging System

The logging system in itself is modular and consists of a number of building blocks. We tried to make the system as general as possible. Thus it is relatively self contained and relies on few assumptions on the logging environment. Figure 1 summarises the

system that is described in the following. Please note that only the Grey components are currently developed.

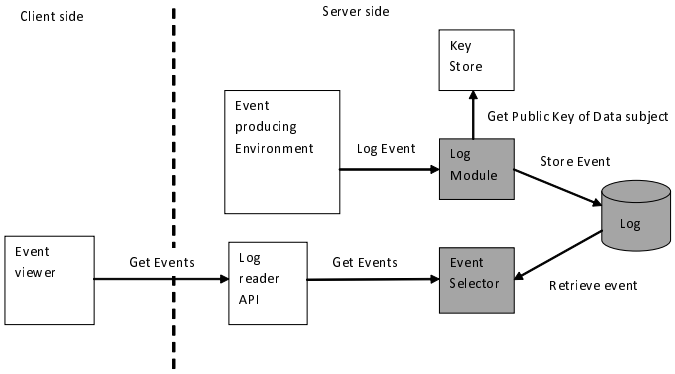

Figure 1: Components in the log system

The event producing environment: This component is under audit and produces log event objects. The event objects consist of at least the subject (i.e., the identifier of the entity performing an action), the action (i.e., what was done), the purpose (i.e., why it was done), the object (i.e., the personal data that the action was performed on), and the data subject (i.e., the identifier of the "owner" of the personal data).[1]

The key store: This server-side protected storage contains public keys of data subjects and hands them out to the log module for a data subject identifier. It depends on the application where the key store is situated and what type of public key used. In the PRIME case the key will be a self signed public key stored together with the PRIME data subject identifier (see footnote 1) in the personal data database on the server side.

The log module: This module receives log events, transforms them into secure privacy preserving log entries with the help of the data subject's public key and stores them in the log.

The event selector: This module, given an entry identifier, retrieves the requested log entry in the log to the requester.

The log reader API: This is a wrapper API that provides controlled or anonymous access to the Event selector depending on the requirements of the service.

The event viewer: This module presents and decrypts the events associated with a data subject in a user friendly fashion. It also contains functionality for searching,

[1]Please note that as soon as personal data is stored in any way in the PRIME system an identifier is generated. This is true even for anonymous access. However, the identifier might not be linked to a known user, i.e., it might be a transaction pseudonym.

sorting, and comparing events as well as for deciding if policy violations have occurred.

Integration in the PRIME core

The integration of the prototype in the PRIME core forced us to implement missing functionality. The missing functionality was primarily to enable two PRIME Cores, one in the client and one in the server role, to exchange secrets and communicate with each other as needed for the transparency logging to function properly. This was accomplished by creating a web-services API and doing some minor modifications to the methods involved when disclosing and storing personally identifiable information. The client, which fetches entries stored in a logging module from a server PRIME Core, was updated to use the API and improved in several ways. However, the client still lacks several features to make it suitable for use by the Data Track, for example persistent storage of entries.

Since the unlinkability properties of the log is dependent on that the log entries cannot be chronologically ordered by an attacker, we were concerned about the underlying database system. Thus we investigated the effects of using HSQLDB[2] as storage on the unlinkability properties of the logging module. Mitigation for some of the flaws found was designed and evaluated, resulting in the conclusion that, while future work is needed, a number of the flaws found could easily be mitigated by configuring HSQLDB properly and randomly distribute the database entries.

Although the work has resulted in a privacy-friendly secure logging module being designed and implemented into the PRIME Core, there are still several areas in which the implementation lacks compared to the design requirements. The client for the logging module still lacks features to be suitable for use by the Data Track. Further research is also needed to make the implementation resistant against the risks posed by memory and disk forensics. We also need to evaluate the performance of the system and effects of possible trade-offs needed in a "real life" setting.

2.2.2 Privacy measurement (Task 2.2.2)

There exist well established models for anonymity focusing on traffic analysis, i. e., analysing properties of single messages as, e. g., timing. However there is only little work done that uses linkability information, that is information about the probability that two messages have been sent by the same user.

Within last year's work, which led to publication [SC09], we have shown how information about linkability between messages (gathered, e. g., by a service provider from knowledge of the content of messages) can be used to reduce sender anonymity beyond what is possible by traffic analysis alone. Therefore, we developed a model where prior knowledge learned from network traffic can be integrated in a "layer-combining model". In this model we denote linkability between messages as a weighted graph.

We have shown lower and upper bounds with regards to the usefulness of linkability information for matching messages to senders. In addition to that, we derived simulation

[2] `http://hsqldb.org/`

results, showing to which extent a matching of messages to senders is possible by using linkability information with different grades of noise.

Model Description

When users communicate with service providers, they may reveal personal information. A service provider can use this information to build user profiles. Some profiles might be linkable with a certain probability, i. e., the service provider can guess that these profiles belong to the same user.

For our model illustrated in Figure 2, we assume a set of users who send their messages to a single service provider, while an anonymity service is obfuscating the relation between senders and messages. The service provider is considered as the attacker, who wants to de-anonymize his users, i. e., he aims at a complete mapping of messages to users. The attacker can distinguish users by observing senders on the network layer and has access to the content of the messages. We further assume that he gains additional information by observing all links in the network, but he is not able to observe the mixing process of the messages. The system is assumed to be *closed*, i. e., all messages are transmitted between nodes within the system, and there are no messages sent to or received from outside the system.

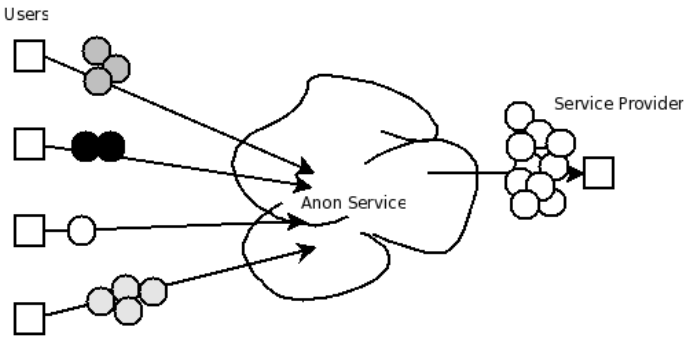

Figure 2: The attacker's view on the system: Users (squares on the left) send messages (circles) to an anonymity service. The service provider on the right hand side receives the anonymized messages and may analyze their content.

With regards to information about linkability between messages on the application layer, we do not explicitly model users by their profiles. We just assume that there exists information about the fact whether pairs of messages have been sent by the same user or not, which, e. g., might be derived from the message contents.

For performing an attack, the attacker has to *cluster* the messages into different groups, where each group corresponds to one sender, i. e., the size of a group of messages corresponding to a given sender is the number of messages sent by this sender. In our work we first derived analytically a lower and an upper bound with regards to the expected number of messages which can be correctly guessed by an attacker. The lower bound is characterized by an attacker we call *random attacker*, who does not have any

linkability information. In this case, any possible clustering having the correct cluster sizes is equally likely for the attacker. The upper bound is characterized by an attacker we call *perfect attacker*, who has complete linkability information, i. e., who knows exactly which messages belong to the same sender, and which do not. Here it is interesting to see that even the perfect attacker may not be able to correctly assign all messages, because among users which sent the same number of messages he can only randomly guess the right assignment of a cluster of messages to a user.

Simulation

Setting. In order to analyze cases, where the quality of the linkability information lays in between the random and the perfect attacker, we performed simulations of such scenarios.

Our experimental setting is as follows: From the knowledge gained from observation of the network layer, the attacker derives the cluster sizes. By analyzing the message content, he derives a weighted graph that represents the knowledge about which messages were probably sent by the same sender. Since we want to abstract from the concrete process of gaining this knowledge by content analysis, we run the attack with the original graph plus noise (circle).

For performing the actual attack, we need to find an optimal clustering, i. e., a clustering where messages that are strongly connected in the graph are assigned to the same clusters. In our simulation, we use *simulated annealing* [KGV83] for finding such optimal clustering. The general idea of simulated annealing is that the algorithm starts with a guessed solution, then randomly picks two elements and swaps these two. If the new solution is better than the old one, it repeats the loop with the newly found solution. Otherwise it continues with a certain probability with either the old or the new solution. The probability that it continues with a worse solution decreases over the running time.

For our problem, we search for the clustering where all messages in the same cluster have been sent by the same sender. Edges between messages from the same sender have more likely a higher weight than others, thus the average of the sums of the edges' weights between messages in the same cluster should be maximal for the clustering where all messages from the same sender are in the same cluster.

Results In Figure 3, a typical result of our attack is shown. 100 messages were sent by 11 senders. On the x-axis the distance between the two Gaussian distributions which were used to add noise is displayed, while on the y-axis the (min, max and average) success rate is displayed. For this example, a random attacker would have a success rate of about 0.1. Note that our attack is already for very small noise distances, i. e., below 1, slightly better. However, for higher noise distances our simulation reaches the theoretical upper bound of 0.81. Furthermore, one can see that the errors are quite large because of the noise impact.

In Figure 4 one can see that the larger the distances between the number of messages different senders have sent are, the faster our attack converges to its individual maximum. Thereby the red pluses represent results from a system where the number of messages sent by each two senders is either equal or differs by at least 5. As we can see, already for small noise distances our simulation reaches its maximum. In contrast to that, the

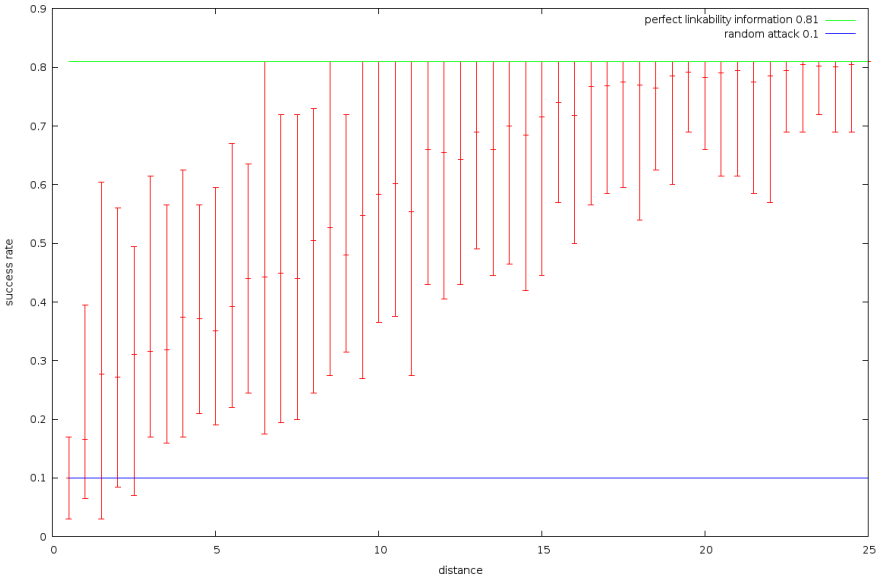

Figure 3: Simulation results for 100 messages, distribution of cardinalities $|c_i|$: [4,5,7,8,9,9,10,10,11,12,15]. For each noise distance displayed, 25 experiments have been made. For each noise distance, the minimum, maximum and average success rate is displayed.

blue stars represent results from a system where the number of messages per sender is much closer to each other. Hence, noise has much more influence on the simulation results since already a small change of the degree of a message might lead to a different clustering.

2.2.3 Privacy-respecting establishment of collaborative groups (Task 2.2.3)

Role Management in Privacy-Enhanced Collaborative Environments

When establishing collaborative groups (an introduction of a privacy-preserving approach was given in [Ste09b, CS09]), determining necessary roles becomes an issue. Applying roles in collaborative groups allows for structuring the group and giving orientation on duties, rights, and expectations related to the group members. So, it is quite natural to approach the management of roles in a privacy-enhanced way while contemplating privacy-preserving establishment of collaborative groups as we did in the second project year [LBP10].

When users work together, each of them takes over a certain position within the group to set up the working scenario. In [Zhu03] it is stated that "without roles, there would be no collaboration". A survey of related scientific literature revealed different interpretations of the concept of roles. According to this, roles can be classified in four main categories:

1. *Positions* (also referred to as *status* or *function*). Roles can be used to describe a

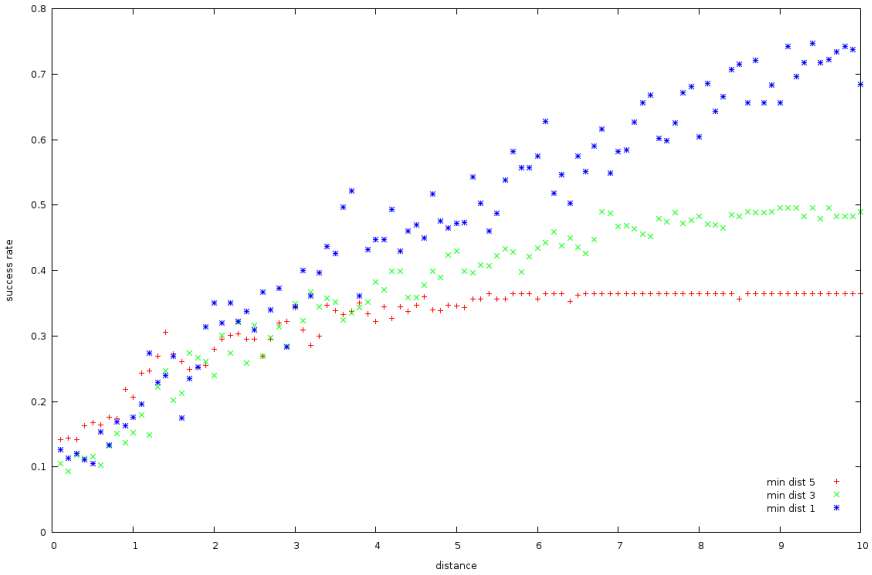

Figure 4: Faster convergence with larger distances between clusters. Red pluses: each two senders sent either the same number of messages or the number of messages differs by at least 5. Green crosses: same number of messages or at least 3 messages difference for each two senders. Blue stars: same number of messages or at least 1 message difference for each two senders

collection of rights, duties [Lin36], and expectations [Luh84]. Often, the terms role and position are used synonymously [Lev66].

2. *Groups.* Roles are also used to categorize users by similarity. In this way, a role displays the kind of user [Zna65].

3. *Behavior.* Roles can be used to assign activities to users [Ger71], e. g., reader or reviewer.

4. *Relations.* Finally, roles can describe different kinds of relationship, cf. [Mea34, CRJ02]. In that case, the role of a user can differ depending on the individual interaction partners, e. g., a secretary is a workmate towards other secretaries, but an employee towards the director.

Collaborative environments require a flexible role management that can be adapted to different situations. Since tasks, authorizations, and team constellations may instantly change during collaborative work, roles have to be adjustable to such conditions. In this context, the following understanding of roles evolved:

> *Roles describe stereotypes of users, which abstract a group of actors with equal rights and duties. Certain expectations are placed in users of specific stereotype addressing the way the users should act like. Further, assignments of roles shall also help the interaction partner to range in a user's position within the collaborative work.*

To develop a highly flexible system that meets the requirements of privacy-preservation, we distinguish the following three dimensions of roles that comply with their management tasks:

1. *Administrative roles* are used to manage users' rights and to realize role-based access control in Privacy-Enhanced Collaborative Environments, e. g., owner or participant;

2. *Functional roles* are used to manage the users' tasks by defining particular privileges, duties, and expectations, e. g., moderator or author;

3. *Group-dynamic roles* are used to identify a user's abilities within a group, e. g., expert or problem solver. To allow for reasonable assignments of group-dynamic roles, the concept of contextual reputation [Laz09] could be used.

In Privacy-Enhanced Collaborative Environments (PECEs), the traditional approach of pre-determined role assignment to user accounts cannot be applied as user accounts as such do not exist. Instead, a policy-credential-based mechanism is used for the purpose of providing authentication and authorization. This is based on certified properties provided by means of anonymous credentials [Cha85] that are issued to the users. Such a credential attests the users particular properties, e. g., the possibility to act under a specific role in given contexts or to execute certain functions in the PECE. To realize access control in PECEs, access control policies (cf. [ADDS05]) are being attached to the resources/services of the PECE. Such a policy indicates which credential a user has

to show in order to get a particular kind of access to the corresponding resource/service. The main advantage of using a policy-credential-based approach of access control is to provide means for advanced specification of controlling access in PECEs, which support collaboration between the users. That way, users specify rules (policies) based on properties other users have to prove, i. e., they do not need to know user names for which they want to control access.

The same is also applied to roles, which the users may indicate instead of user names, denoting users as entities acting under particular roles. In comparison to the well-known role-based access control mechanism, this approach does not require a centrally managed list of users, which are assigned to a particular role. Instead, privacy is supported by externalizing that list and decentralizing the user-role assignment. That way, users may switch their partial identities according to the corresponding application context and they may decide if they change or retain the roles displayed to the other users.

In conclusion, integrating roles in collaborative environments helps to maintain the focus of the tasks. In PECEs, this point gains even more importance when users frequently switch between different partial identities. Role profiles and descriptions can remind the users of PECEs of their aims, duties, or relationships. Thereby, role management can be provided in a privacy-enhanced way by using the particular feature of PECEs to support partial identities. This means, different roles assigned to one and the same person may be distributed to different partial identities of that person.

However, the assignment of roles implies adding information to partial identities. This may potentially (in dependence of the size of the resulting anonymity set) help to increase the linkability of activities of that user or, in the worst case, it may even result in identifying the user unintended by her/him. In result, work with roles in (privacy-enhanced) collaborative environments is as privacy-preserving as the user partitions his personal data and selects the right partial identity for disclosure.

Privacy-Enhanced Event Scheduling

Event schedulers, used for the cooperative identification of the best time for an event involving many users and well-known from social software or stand-alone Web 2.0 applications [Näf10], typically share the problem that they disclose availability patterns of their users. Privacy-enhanced event scheduling can be understood as an instance of an electronic voting scheme. However, the main difference between existing e-voting schemes and the event scheduling problem is the parameter in which the system has to scale efficiently. Typical governmental or committee voting schemes in the literature scale in the number of *voters*, each of whom has one (or a few) vote(s). By contrast, event scheduling requires a limited number of voters to make many binary choices, one for every time slot. Hence it has to scale in the number of *independent choices*. As opposed to typical governmental election schemes, which have to cope with millions of voters, it is not so crucial for event scheduling to scale gently in the number of participants. Typically we schedule events only for small closed groups.

Derived from typical requirements in e-voting, the following requirements apply to privacy-enhanced event scheduling:

Verifiability Every voter should be able to verify that no other voter has cheated and that her own vote has been counted.

Privacy Nobody should learn more than absolutely necessary about the availability of other voters and thus should not be able to infer on their identity, i. e., every participant should only learn that the other participants are available at the one specific date which was chosen.

Untrusted server As little trust as possible should be placed in any central entity (e. g., the vote server).

Usability The scheme should not require many more steps than existing event schedulers, i. e., it should not require much more user interaction and message exchanges.

Efficiency The scheme should be more efficient than existing schemes. More precisely, it should be efficient for large scheduling problems with many possible event dates.

Our proposed scheme consists of three mandatory phases and one optional verification phase that is run when inconsistencies occur. In the following we describe the phases in more detail.

Poll Initialization In the poll initialization phase, the initiator has to define the set T of all possible time slots at which the event can take place. For now we consider a closed group, so the initiator has to define the set of voters P.

To ensure the anonymity of each voter, we propose to use superposed sending, generalized to other abelian groups than GF(2) [Cha88]. This anonymity scheme has an implied homomorphism, in which we can integrate the voting protocol. Let n be the modulus that defines the group $\mathbb{Z}_n$ for superposed sending. To ensure that no overflow occurs while summing up all votes, the modulus n has to be large enough, i. e., $n > |P|$.

All voters have to exchange symmetric keys beforehand. To be precise, each voter has to do a key exchange with $|P| - 1$ other voters. Therefore he exchanges with each other voter, $|T|$ independent and uniformly distributed random numbers $r_t \in \mathbb{Z}_n$. This means he ends up with $(|P| - 1) \cdot |T|$ random numbers. To detect cheating by key modification later on, the participants have to sign their keys. A digital signature by participant p of message m will be denoted as $\mathrm{sig}_p(m)$. Keys shared by a pair of voters are denoted as $k_{p_i,p_j,t}$ where $p_i, p_j \in P$ and $t \in T$. To run the key exchange, each pair of voters p_i, p_j where $i < j$ does:

1. Exchange a random number $r_t \in \mathbb{Z}_n$ for every $t \in T$.
2. Voter p_i stores the signatures $\mathrm{sig}_{p_j}(k_{p_i,p_j,t})$ and the keys $k_{p_i,p_j,t} = \quad r_t$.
3. Voter p_j stores the signatures $\mathrm{sig}_{p_i}(k_{p_j,p_i,t})$ and the keys $k_{p_j,p_i,t} = -r_t \bmod n$.

Afterwards each voter p_i holds a key matrix $\boldsymbol{k}_{p_i}$. Rows in this matrix contain the keys and signatures exchanged between two voters p_i, p_j and are denoted by $\vec{k}_{p_i,p_j}$.

Casting of Votes In this phase, each voter has to state for every given time slot whether he or she can participate in the event or not. The superposed sending, as underlying anonymity mechanism, sums up all messages. This suits well with our goal, as we are only interested in the sum of the given votes.

Each voter p has to calculate an encrypted vote vector $\vec{d_p}$ which consists of $|T|$ encrypted votes $d_{p,t}$,

$$\vec{d_p} = (d_{p,t_1}, d_{p,t_2}, \ldots, d_{p,t_{|T|}}) \,. \tag{2.1}$$

An encrypted vote $d_{p,t}$ is calculated by adding all keys the voter has exchanged with other voters to the actual vote $v_{p,t}$ modulo n,

$$d_{p,t} = v_{p,t} + \sum_{i=1, p_i \neq p}^{|P|} k_{p,p_i,t} \bmod n \,. \tag{2.2}$$

$v_{p,t} \in \{0, 1\}$ is the specific vote with the semantic that value 0 means the voter p is unavailable at time slot t and value 1 signals availability.

Vector $\vec{d_p}$ is sent to the server and will be published there after all voters have casted their votes.

Result Publication When all vote vectors are published on a central server, any party can calculate the sum of these vectors. Since all keys should be pairwise inverse, the result should be a vector $\vec{v} \in (\mathbb{Z}_{|P|})^{|T|}$ containing the sum of all votes at the specific time slots.

As *selection rule* for the agreed time slot t_a, the earliest time slot for which $t_a = |P|$ is chosen.

Result Verification A malicious voter could send values $v_{p,t}$ different from 0 or 1. Since this would be completely invisible to the others, the result has to be verified after publication. The first step is that everybody checks if he voted for this specific time slot. This is done in most cases anyway, when the voter inserts the event in his personal calendar. If a voter discovers an inconsistency here, he can request the decryption of the individual votes for the agreed time slot t_a. If everybody stuck to the protocol, this is no privacy problem because one can infer that everybody must have sent a 1 from the mere fact that the time slot has been selected. This means, after t_a is found, every voter p_i has to publish his signed shared secret keys $k_{p_i,p_j,t_a}, \mathrm{sig}_{p_j}(k_{p_i,p_j,t_a})$ for the selected time slot if at least one voter demands it. With these keys, the respective elements d_{p_i,t_a} of the encrypted vote vectors $\vec{d_{p_i}}$ can be decrypted and it can be verified that every voter casted a 1. However, a malicious voter can send values higher than 1 in an attempt to compromise the privacy for a specific time slot.

Conclusion We proposed a privacy-enhanced event scheduling scheme that is suitable to be implemented in practical Web 2.0 sites, or groupware applications. A more detailed description of the protocol with several enhancements and a detailed security analysis was given within two other academic papers [KB09, Kel09]. In addition, a partial implementation of the scheme has been done [Kel10]. The scheme is efficient and scales, in terms of hash and symmetric encryption functions, linearly in the number of possible points in time. The effort per voter in terms of asymmetric cryptographic operations scales linearly in the number of voters. Moreover, no central trusted entity is required.

2.2.4 Trust management by interoperable reputation systems (Task 2.2.4)

Reputation systems need to help users in interactions to estimate their interaction partners' behaviour beforehand, but they also need to respect privacy of all users involved. While in the first project year we focused on guaranteeing rater anonymity by using convertible credentials [Ste09a], in the second project year we concentrated on using anonymous payment systems as an alternative approach.

There are already two proposals of system designs following this approach [ACBM08, Vos04]. A drawback of both systems is that the *liveliness of reputation* cannot be guaranteed. Liveliness of reputation means that reputation should always consider all recent interactions or give users an indication that there are no more. Especially the reputation system should not offer users the possibility to reach a final state in which bad behavior no longer damages their reputation.

System environment

For our system environment, shown in Figure 5, we assume a community system supporting pseudonymous interactions among users. This might be, e. g., a marketplace such as eBay where every user might be a seller (provider) or buyer (client). Let M be such a user offering interactions under the pseudonym P_M to other users. The community deploys a reputation system provided by a reputation provider ReP. The reputation system collects positive and negative experiences of users' behavior during interactions. Thus we assume that only interaction-derived reputation is aggregated by our system. If a user U becomes interested in the interaction offered by P_M, U inquires P_M's reputation under pseudonym P_{U_1}. If U decides to take part in this interaction, U uses another pseudonym P_{U_2} to interact. Afterwards, U rates P_M using a new pseudonym P_{U_3}. M can now include the rating P_M got in the overall reputation account at ReP.

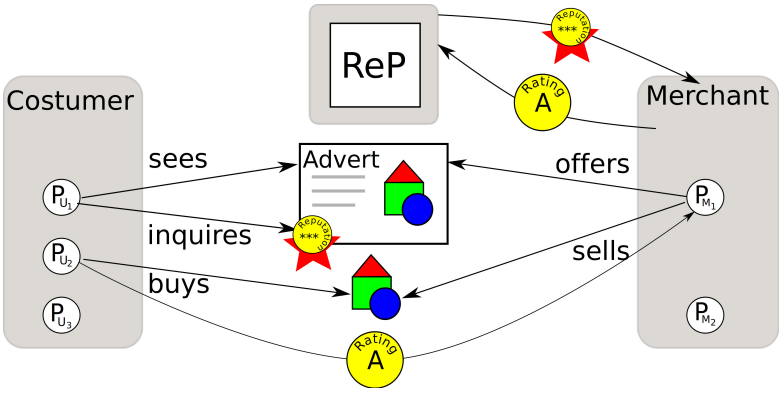

Figure 5: Model of system environment

The requirements a privacy-respecting reputation system has to fulfill are derived in [Ste09a, ENI07]. For our system environment we distinguish two types of attackers, namely, the privacy attacker and the security attacker.

Privacy attacker. As privacy attacks we subsume attacks on *raters', inquirers' and ratees' anonymity.* We assume that reputation can be queried anonymously (e. g. by its publication on a website, as it is the case for eBay) and therefore we concentrate on raters' and ratees' anonymity. We assume that the privacy attacker cannot observe who is communicating with whom, that is, all users are communicating via an anonymity service. Furthermore, the attacker might collude with the reputation provider, but cannot cheat on the reputation values, that is, he is an honest but curious attacker. In addition, the privacy attacker can only control a limited number of users so that a sufficient large anonymity set (which contains the users not controlled by the attacker) is preserved.

Security attacker. We see the security attacker as an attacker on the *integrity and authorizability of ratings* and on the *fairness and liveliness of reputation.* We assume a global attacker who might observe all interactions, but that cannot control the reputation provider. We show in our analysis that an attacker that controls all users in the system can only forge a reputation if the attacker can break the underlying eCash system or forge the credential representing the reputation itself.

System design

The reputation provider ReP keeps an interaction account and a reputation account for every user. ReP thereby guarantees that every interaction is actually rated, possibly also in a negative way, and considered for the user's reputation. We implement both accounts as accounts of an anonymous payment system and the ratings and interactions both as coins. Thereby, negative coins can be implemented by two instances of an anonymous payment protocol with a joint account, where coins of the first system are counted as $+1$ and coins of the second one as -1. For this, we use two instances of the protocol from [ACBM08].

- *Interaction counter:* This instance is used to count the number of interactions a user U was involved in and should be rated for.
- *Reputation counter:* The other instance aggregates the ratings received, both positive and negative ones.

The protocol is given in the following and summarised in Figure 6.

Registration. In order to initialize the reputation system, every user withdraws a wallet from ReP, which contains n interaction coins (S_i, π_{it}) and reputation coins (S_i, π_{ir+}) for positive ratings and (S_i, π_{ir-}) for negative ratings. The coins are issued in triples with the same serial number S_i and π_{it}, π_{ir+} and π_{ir-} are the double-spending tacks and signatures with $i = 1 \ldots n$.

Interaction. If user U wants to interact with an interaction partner M (whom U knows as P_M) using a pseudonym P_U, U starts the interaction by awarding an interaction coin (S, π_t) to P_M. P_M spends this coin to its registered pseudonym M, which deposits this coin and requests a one-show credential from the reputation provider stating the fact that the number of coins in the interaction account has been increased. P_M shows this

credential to P_U. Now the actual interaction can take place. Furthermore every party needs to check the age of the coin to prevent undetectable double spending, as outlined in the analysis below.

Rating. After an interaction, P_U rates P_M by awarding a reputation coin (S, π_{r+}) or (S, π_{r-}). P_M deposits this coin. During the deposit the reputation provider checks whether the serial number of an earlier deposited interaction coin equals the serial number of the reputation coin to avoid that M uses one of his own coins to rate himself with a positive rating instead of the (possibly negative) one received from P_U. As for the interaction coins, the age of the rating coins needs to be consistent.

Showing Reputation. If users want to show their reputation to someone, they need to request an up-to-date reputation credential with a time stamp from the reputation provider. The reputation provider issues a reputation credential only if the interaction account and the reputation account contain the same number of coins or they only have a small difference for highly active users. The reputation provider can also play the role of a global *time provider* in a very natural way by using the number of total (by every user) deposited coins as global time. This also gives an estimate on how much users could cheat about their reputation, since the time difference between issuing the credential and now is the maximum number of possibly negative coins a user could have received in between.

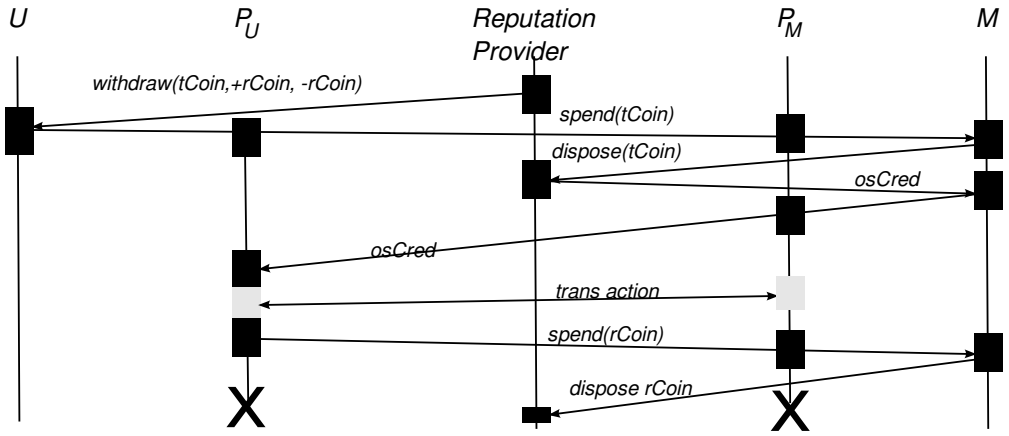

Figure 6: The reputation granting is bound to interactions.

Batching. The protocol presented above might still raise timing issues on users' anonymity. In order to minimize this problem, we propose to batch all user activities in rounds of three phases. In every round users get wallets with n coin triples and a sufficient amount of credentials about their reputation level, which they achieved in the round before. After that, users find their at most n interaction partners (using the credentials) and spend on them an interaction coin. In a second phase the interaction

partners deposit their interaction coins and the actual interaction takes place. After the interaction the users spend on their interaction partner a reputation coin with the intended value. In the third phase all interaction partners deposit their reputation coins.

Privacy and Security Analysis

Security Attacker. The interaction registration phase depends on the security of the transferable eCash system: even if all users collude, a double spending can be proven and traced back to its origin. The user U, who starts the interaction, cannot forge the interaction coin without revealing his registered user name U, since the dispose algorithm would recognize this double spending. The user M, however, might transfer the coin multiple times from P_M to M. In this case the deposit algorithm will return a proof that P_M double-spent the coin, where P_M is a non-registered pseudonym. However, since the number of hops for a coin is known, only a pseudonym controlled by M can double spend. Since M needs to reveal its identity to the reputation provider, it can get its deserved punishment in case of double spending. The argumentation for the rating is similar. These properties ensure the security properties *integrity and authorisability of ratings*, as well as *fairness of reputation*.

Privacy Attacker. *Anonymity of the inquirer* can be guaranteed by inquiring with a one-time pseudonym or publication of P_M's reputation. The *rater's anonymity* against the reputation provider is perfectly preserved by the anonymous payment system: the rater is anonymous among all the users who withdrew interaction and reputation coins during this round. The *ratee's anonymity* M cannot be guaranteed because the disposal of the interaction coin before the interaction and the reputation coin after the interaction are in principle linkable to M. This is not a problem as long as it is assumed that ReP cannot observe any peer to peer traffic. Batching allows relaxing this condition. Assume that ReP can observe which peers communicate, then ReP could link a P_U with its corresponding U if there is only one user who deposits a coin at this time. If it is assumed that many users deposit their coins at the same time, these users would be anonymous among each other. Batching allows concentrating these steps. Furthermore, batching helps to protect naive users from outside attackers who re-query the reputation of their interaction partners, since in every round the reputation of a user stays constant. However, batching is the more effective the longer the rounds are, but the longer a round is the longer a malicious node stays unpunished. The right trade-off between security and privacy depends on the application and is beyond the scope of this work.

2.2.5 Privacy awareness (Task 2.2.5)

Besides working on the theoretical background [Ste09b, CS09] and developing a first version of a tool to support privacy-awareness [Ste09c], we also did empirical research in order to study the effect of privacy-awareness information (in short: PAI) on end users. In this section we explain the concept of this study and present the main results.

Hypotheses

In order to research the effect of PAI on the users' perceived level of privacy and on their communication behavior in a web forum we did an empirical study [Pöt09b]. In this study we used the privacy-awareness tool described in [Ste09c]. We based the hypotheses for our study on the theory of the cues-filtered-out approach (CFOA). The CFOA implies that, due to missing contextual cues in computer-mediated communication (CMC), individuals are more lavish regarding the disclosure of personal data and that they behave more unsocial than in face-to-face settings [SK86, Her02, Dör08]. Transferred to the behavior of forum users, this means they tell a lot about themselves on the Internet without considering the potential audience and also tend to insult others, e. g., the phenomena of "flaming" in forums [Lee05]. We investigated whether the display of PAI as additional cues in a forum helps to reduce the effects predicted by in the CFOA. Our concrete hypotheses were as follows:

> Subjects who are provided with privacy-awareness information
> **H1**: feel they have less privacy during their work with the forum,
> **H2**: disclose less personal data in their postings, and
> **H3**: post less off-topic statements
> than the participants from the control group.

Study Design

Our study consisted of two main parts:

Part A. In the practical part, the subjects visited a forum, which was carefully prepared by us including all the postings that were available when a participant entered the forum.

Part B. In a following on-line questionnaire we asked subjects about their perceived level of privacy during the work with the forum and we also collected demographic data.

We published invitations to participate in the study on mailing lists and forums on the Internet. To avoid bias, we used a cover story and told that we were doing research on participants' interest in different topics and their participation in Web 2.0 communities.

Table 2: 2 × 2 design of the study

	Audience shown?	
IP shown?	yes	no
yes	ExG_1	ExG_2
no	ExG_3	CG

The forum in part A was a 2 by 2 design that distinguished two types of PAI as independent variables: the display of the potential audience and the presentation of the subjects' IP address (Table 2). Participants of the study were randomly assigned either to one of

the three experimental groups (ExG_i) or to the control group (CG). The participants in the three experimental groups were shown the respective PAI in an orange box placed on top of the forum (see Figure 7). The IP address was intended for reminding the subjects that they are not as anonymous as they might feel, whereas the indication of the potential audience should compensate partly the missing context cues about communication partners. Besides the display of the PAI, there was no further difference in the appearance of the forum for the different groups.

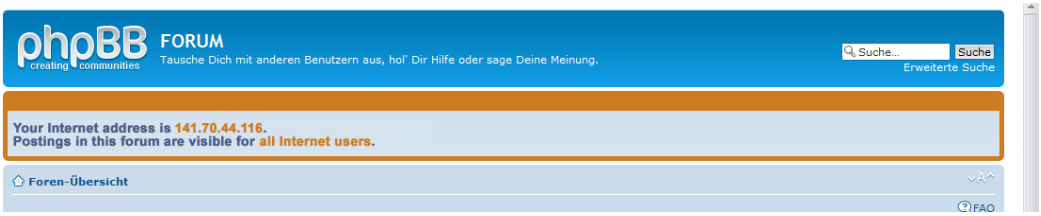

Figure 7: User interface of the forum for the experimental group ExG_1 with IP address and potential audience as privacy-awareness information shown in the orange box (originally shown in German).

Data analysis and results

156 participants visited the forum (part A) at least as lurkers and completed the questionnaire (part B). 97 of them wrote altogether 330 postings in the forum, which means that 38 % of the subjects decided not to contribute to the forum, but to answer the questionnaire. After the field study phase has ended, two trained coders rated all postings from the forum in part A according to a five point rating-scale that we developed as an enhancement of a similar instrument used in [BGO07]. Our scale permits to evaluate each forum posting i (i=1..n) of a subject s in the five subcategories *personal information* (PI), *personal thoughts* (PT), *personal feelings* (PF), *content divergence* (CD) and *insulting statements* (IS) [Pöt09a]. Using the average rating scores of the two coders per posting, we further calculated the values for the disclosure of *personal data* (PD) and *off-topic statements* (OTS) for each subject (cf. Equations 2.3 and 2.4) and did an analysis of variance (ANOVA) to test our hypotheses. Finally, we calculated the *Perceived Privacy Index* (PPX)[3] from the participants' answers in part B to the question how public, how private and how anonymous they have felt during their visit of the forum. Each item was measured on a 0 to 100 % slider scale. The higher the PPX value, the more private a subject has felt.

$$PD_s = \frac{\sum_{i=1}^{n} PI_{s_i} + PT_{s_i} + PF_{s_i}}{n} \quad (2.3) \qquad OTS_s = \frac{\sum_{i=1}^{n} CD_{s_i} + IS_{s_i}}{n} \quad (2.4)$$

[3]Part B contained the question "Please indicate to which extent the following adjectives describe your feelings while working with the forum: 0 % (not at all) – [*adjective*]– 100 % (very much)?" (originally asked in German). The PPX is composed of the adjectives "public" (scale inverted), "private", "anonymous".

Hypothesis H1 was supported by the results of the statistical analysis, which prove that the PPX was significantly higher for participants in the control group than for those in the experimental groups (Table 3). The data did not confirm the hypotheses H2 and H3. Providing PAI neither decreased the disclosure of PD (mean of 1.92 for all ExG_i vs. 1.81 for CG), nor significantly influenced OTS (mean of 0.64 for all ExG_i vs. 0.72 for CG). An explanation for these null-results may be the fact that some participants did not clearly notice or understand the PAI. Besides, due to the experimental setting, participants visited the forum only once. Hence, there were no established relationships to other forum users and that we could not gather and analyse contributions over a longer period in time.

Table 3: Perceived Privacy Index for different groups according to the presented privacy-awareness information

	PPX				
	Mean	Min	Max	n	p
All	121.64	0	290	156	
No PAI shown	132.03	0	290	71	0.05
others	112.96	10	268	85	
IP or both shown	110.00	10	268	60	0.05
others	128.92	0	290	96	
Aud or both shown	116.52	30	268	50	0.46
others	124.06	0	290	106	
Both shown	112.96	30	268	25	0.43
others	123.30	0	290	131	

p from ANOVA (F-test)

2.3 Future research

2.3.1 Transparency support tools (Task 2.2.1)

Further research in transparency support tools will focus on three main issues. First we would like to more thoroughly evaluate the performance and privacy compliance of the log. Secondly we will investigate how cascading logs can be handled. This is something that can happen if, e.g., the data about a data subject is released to a third party. In this case there would be a log in the data controller's system and in the third party system without the data subject taking part in the transfer. Last, we will investigate how the log information can be incorporated in the PRIME data track and how this information should be presented to the user. The last issue will be addressed in cooperation with Activity 4.

2.3.2 Privacy measurement (Task 2.2.2)

Our future research will elaborate mainly on two issues. The first one will be to enhance the work of this year by improving the underlying system model as well as the methods for anonymity analysis. Therefore, we will extend our model to more comprehensive network as well as application layer models. A second issue will be further research on how we can use our model to directly derive sender anonymity measures in terms of Shannon entropy as metric, that is calculating the probability distribution that a given user has sent a given message. Further, we will go into more detail with respect to methods and measures for quantification of anonymity and their usability for different scenarios.

2.3.3 Privacy-respecting establishment of collaborative groups (Task 2.2.3)

In the near future, we plan to implement the missing features into our Web 2.0 application [Kel10]. Future extensions could complement the features, e. g., by integrating a threshold scheme, dynamic insertion and deletion of time slots, updating and revoking votes. The risk of denial-of-service or legal-but-selfish votes could be reduced by letting voters prove that they signaled availability for more than a certain minimum number of time slots. Additionally it may be desirable to have more complex decision rules than just maximum agreement.

2.3.4 Trust management by interoperable reputation systems (Task 2.2.4)

Although the analysis of our reputation system based on anonymous payment is already quite promising for actual deployment, our future theoretical research will concentrate on denial of service prevention and on the privacy problems caused by traffic analysis. Furthermore, the problem of self rating needs to be solved. For this reason we concentrate our practical work on tools on the Jason tool we presented in [Ste09b, CS09, Ste09c] and the implementation of a reputation system as the non-focal prototype of WP1.1.

2.3.5 Privacy awareness (Task 2.2.5)

In our study, we showed with empirical evidence that the display of PAI effects the perceived level of privacy of forum users. This implies, that the presentation of PAI is a reasonable approach to support forum users in making informed decisions about the disclosure of personal data. A next step will be to improve the privacy-awareness tool by finding a representation of the PAI that is even better suited to make users aware of their actual level of privacy. We also will test users' understanding of different kinds of PAI, e. g., showing of IP address vs geolocation.

Chapter 3

Privacy of data (WP 2.3)

3.1 Introduction

Data holders find it increasingly difficult to produce anonymous information in today's globally interconnected society. The technology advancements leave the information vulnerable to inference and linking attacks, meaning that from data that seem anonymous and from available public data, an adversary can make some inferences about sensitive data. This is possible because released information often contains other data that in combination can be linked to publicly available information to re-identify users. Existing approaches that are based on data distortion or disturbance do not satisfy the requirements of novel scenarios, where data must be made available subject to the constraint that data themselves must be truthful while preserving the privacy. In fact, data quality is of critical concern to organizations since data and information become key strategic resources and therefore poor quality can have significant consequences on the ability of organizations to act effectively.

The focus of Work Package 2.3 is then represented by the consideration of large data collections that contain sensitive information on citizens. The overall goal is the definition of novel metrics and techniques able to support the management of privacy requirements, at the same time offering a significant degree of utility in access to the data. The investigation of these topics in PrimeLife has two roles: on one hand, it can produce concrete techniques for the protection of personal information, identifying the amount of exposure deriving from access to the data and proposing approches able to satisfy the stringent privacy requirements that the project wants to support; on the other hand, the identification of metrics and techniques can provide input on the definition of components in the policy language able to express user preferences on the processing of their data (e.g. users can express the preference that their data when collected do not have to expose the user profile in groups with cardinality less than k). The availability of relational database technology is assumed for many of the scenarios considered in the Work Package. Work Package 2.3 is organized in the following three tasks, focusing on specific aspects of the privacy problem.

- *Task 2.3.1 Privacy assessment and privacy metrics* will define metrics and measures for formally describing the level of privacy protection that is guaranteed on a data collection.
- *Task 2.3.2 Techniques for enforcing data privacy* will focus on the definition of techniques for protecting data privacy according to some constraints that must be satisfied by the data themselves.
- *Task 2.3.3 Efficient organization and access to privacy-preserving data collections* will investigate the efficiency aspects related to the techniques enforcing data privacy, which are fundamental aspects to take into consideration to guarantee the applicability of the data protection techniques in very large data collections.

In the second year of the project, Work Package 2.3 has continued the work toward the developement of novel solutions for enforcing privacy constraints in mobile networks and privacy requirements/constraints within business applications. Some effort has been dedicated to the construction of a prototype supporting the definition of confidentiality constraints on a relational schema and the identification of a fragmentation satisfying them. In the following, we briefly describe these research results and highlight a number of possible directions as future work.

3.2 Research results

3.2.1 Privacy assessment and privacy metrics (Task 2.3.1)

Support tool for anonymization

Data anonymization is a common used method for removing private information before releasing data. However, finding the right balance between privacy and the quality of data is often difficult, and it needs a fine calibration of the anonymization process. It includes choosing the 'best' set of masking algorithms and an estimation of the risk in releasing the data. Both these processes are rather complex, especially for non-expert users. We developed a tool [BMST09, TSBM09] for assisting the user in the choice of the set of masking transformations. We also proposed a caching system to speed up this process over multiple runs on similar datasets. This tool analyzes a given disclosure policy, then estimates the privacy risk for a given dataset (compared to a target risk value) and, in case, it provides the user with relevant suggestions (fields to be additionally masked) to decrease the risk up to the set threshold. The tool implemented a caching system to optimize this process (which in large datasets may be extremely long, therefore, limiting the applicability).

Privacy metrics

Regarding Privacy metrics, we investigated how different metrics (such as k-anonymity, l-diversity, t-closeness) may be expressed in a common framework. These measures are able to capture different aspects of privacy risk, in particular k-anonymity relates to *identity disclosure* risk, whereas l-diversity and t-closeness (as well as δ-privacy) measure different features of *attribute disclosure* risk. Ideally, we should be able to express these metrics

in terms of semantically "similar" measures, so we can combine them in a optimization problem. To this scope we proposed an information theoretic formulation for the privacy metrics [Bez10]. Current information theory approaches are based on the usage of mutual information, but since mutual information is an average quantity, it is not completely able to capture the risk at the level of single records. Thus, we introduced one-symbol information and expressed the three metrics as contribution to information. We are currently exploring how this framework can be used to formally describe an optimization problem between privacy and utility in real-life scenarios.

Protection of statistical information

We studied different inference based attack models for public security systems (e.g., Human Resource system for defense industry), with the goal to prevent that confidential statistical information may leak from these systems. An example scenario is the following: A military organization holds a large amount of sensitive data, and keeps them in a very secure environment (referred to as High). They include strategic documents, plans, etc., but also personal information about the army personnel, such as soldier health records. They may need to transfer some of these data to a less secure environment (referred to as Low), for example, they have to release the health records to hospitals or to the data subjects. These data may be sanitized (for example, removing information about the exact location where a certain accident happened) before release. Still, there is the risk that an unlegitimate user could access, possibly breaking the security in the Low environment, these data, and aggregating them to infer confidential information, stored in the High environment. For example, aggregating the personal data records from a large group of soldiers may derive that in a certain location most of the soldiers have a certain disease or simply they are older than average, indicating that a site hosts mostly officials or the reserve force. In this context, the problem is how to estimate the risk of inference of statistical information. Compared to the typical scenarios of statistical disclosure, where the personal information contained in the microdata have to be protected, here, we want to avoid the inference of statistical quantities from the microdata. To this scope, we proposed a risk model based on information theoretic measures, which permits to estimate how much information has been released compared to a baseline (typically the publicly available data). We are currently developing a prototype for supporting the user in estimating the risk before release, and, possibly applying mitigation measures (e.g., data masking).

3.2.2 Techniques for enforcing data privacy (Task 2.3.2)

The work on the design of techniques for data privacy has mainly investigated the management of privacy constraints in relational databases.

Privacy Constraints in Relational Databases

The design of distributed databases and associated techniques have been a topic of interest in the 1970s, at the birth of relational technology. The scenario that was considered in those years was significantly different from the current scenario: the emphasis was on the implementation of systems owned by a large organization managing information

systems in several centers that offered the opportunity for the processing of distributed queries. The current ICT scenario is instead characterized by many important opportunities for the use of distributed databases where previous assumptions do not hold. Two important differences compared to traditional approaches are: *(1)* the need to integrate the services of database providers that do not belong to the same organization; *(2)* the presence of a variety of platforms, with an increase in the number and availability of devices that have access to a network connection, together with the presence of powerful servers offering significant computational and storage resources. The first aspect forces the requirement to specify security functions limiting access to the information stored in the databases. The second aspect instead forces an environment where the data and computational tasks are carefully balanced between the lightweight device and a powerful remote server. The two aspects are strictly related, since the servers are typically owned by service providers offering levels of cost, availability, reliability and flexibility difficult to obtain from in-house operations. Note that any device is classified as "lightweight" that exhibits performance or storage features significantly worse than what can be offered, for the specific application, by a remote service provider. Mobile devices certainly fit this description, but the scenario can be extended to generic computational platforms.

The motivation of this work lies in the desire to define novel solutions for the processing of large data collections, which are assumed managed by traditional database technology, in a scenario where there is interest in using the services of a honest-but-curious powerful server, with a robust guarantee that confidentiality of information is protected.

In the literature, this problem has been addressed by combining fragmentation and encryption, thus splitting sensitive information among two or more servers and encrypting information whenever necessary [ABG+05, CDF+07, CDF+09b]. In [CDF+07, CDF+09b] a relation r is split into *two or more fragments possibly stored on one server* and encryption is used to protect single sensitive attributes. Furthermore, for query execution efficiency, attributes that are not represented in the clear within a fragment are represented in encrypted form, providing the nice property that each fragment completely represents the original relation. The consequence of this design choice is that to evaluate a query, it is sufficient to access a single fragment, thus avoiding join operations (needed with the previous paradigm), which are quite expensive. This solution however still requires the client to possibly decrypt the attributes appearing in encrypted form in the fragment for evaluating a condition on them or for returning them to the user.

A common assumption of these solutions is that encryption is an unavoidable price to be paid to protect the information. There are, however, situations where encryption may not be applicable for protecting information. One issue can be the computational load imposed by encryption. For instance, in systems more constrained in battery power rather than memory capacity, it is beneficial to spend some memory space to save on power consumption. Also, in real systems keys can become compromized, and keys can become lost, making the protection of the system dependent on the quality of key management services, rather than on the quality and strength of the cryptographic functions. Since key management is known to be a difficult task, an encryption-less solution can be of interest for many important applications.

To address these situations, a paradigm shift is proposed, where information is protected without encryption. The basic idea is that a small portion of the sensitive in-

formation can be stored on the client, trusted for both managing the information and for releasing such a sensitive information only to the authorized users, while the remaining information can be stored on an external server. Obviously, from the information stored on the external server it should not be possible to reconstruct a sensitive association (confidentiality constraint) since otherwise a privacy violation occurs. Since the assumption that the external servers are all honest-but-curious is still valid, the client is the only entity in the system that can manage a portion of sensitive data. Sensitive data and associations can then be protected by splitting the original relation r into two fragments, denoted F_o and F_s, stored at the client and at a honest-but-curious storage server, respectively.

The syntax for the expression of the confidentiality constraints is the one designed in this task in the first year of the project. We describe here the model that underlies all the works in this area. We consider a scenario where, consistently with other proposals (e.g., [ABG+05, CDF+07]), the data to be protected are represented with a single relation r over a relation schema $R(a_1,\ldots,a_n)$. We use the standard notations of the relational database model. Also, when clear from the context, we will use R to denote either the relation schema R or the set of attributes in R.

Privacy requirements are represented by *confidentiality constraints*, which express restrictions on the single or joint visibility (association) of attributes in R. Confidentiality constraints are formally defined as follows [ABG+05, CDF+07].

Confidentiality constraint: Let $R(a_1, \ldots, a_n)$ be a relation schema, a *confidentiality constraint* c over R is a subset of the attributes in R.

While simple in its definition, the confidentiality constraint construct supports the definition of different privacy requirements that may need to be expressed. A singleton constraint states that the *values* assumed by an attribute are considered sensitive and therefore cannot be accessed by an external party. A non-singleton constraint states that the *association* among values of given attributes is sensitive and therefore should not be outsourced to an external party.

Figure 8 illustrates plaintext relation Patient (a) and the confidentiality constraints defined over it (b).

- c_0 is a singleton constraint indicating that the list of SSNs of patients is considered sensitive.

- c_1 and c_2 state that the association of patients' names with their illnesses and with the causes of the illnesses, respectively, is considered sensitive.

- c_3 and c_4 state that the association of patients' dates of birth and ZIP codes with their illnesses and with the causes of the illnesses, respectively, is considered sensitive; these constraints derive from c_1 and c_2 and from the fact that `DoB` and `ZIP` together could be exploited to infer the name of patients (i.e., they can work as a quasi-identifier).

- c_5 and c_6 state that the association of patients' jobs with their illnesses and causes of the illnesses, respectively, is considered sensitive, since it could be exploited to establish a correlation between the job and the illness of a patient.

PATIENT

SSN	Name	DoB	ZIP	Job	Illness	Cause
123-45-6789	Alice	80/02/11	20051	Secretary	asthma	Dust allergy
987-65-4321	Bob	85/05/22	22034	Student	fracture	Car accident
147-85-2369	Carol	73/07/30	22039	Secretary	carpal tunnel	Secretary
963-85-2741	David	75/11/26	20051	Lawyer	hypertension	Stress
789-65-4123	Emma	90/03/15	22035	Student	asthma	Student
123-65-4789	Fred	68/08/07	22034	Accountant	hypertension	Wrong diet

(a)

$c_0 =$ {SSN}
$c_1 =$ {Name,Illness}
$c_2 =$ {Name,Cause}
$c_3 =$ {DoB,ZIP,Illness}
$c_4 =$ {DoB,ZIP,Cause}
$c_5 =$ {Job,Illness}
$c_6 =$ {Job,Cause}

(b)

Figure 8: An example of relation (a) and of confidentiality constraints over it (b)

The satisfaction of a constraint c_i clearly implies the satisfaction of any constraint c_j such that $c_i \subseteq c_j$. We are therefore interested in enforcing a set $\mathcal{C} = \{c_1, \ldots, c_m\}$ of *well defined* constraints, where $\forall c_i, c_j \in \mathcal{C}, i \neq j,\ c_i \not\subset c_j$.

The publications of the second year [CDF+09a, CDF+09c] start form this model and describe how the above confidentiality constraints can be supported in a scenario with a client and a server, without using encryption. In [CDF+09a] the model is presented for the first time, providing motivation and the basic features of the model. In [CDF+09c] the approach is extended, describing a heuristic technique for the identification of an encryption-free fragmentation between client and server able to support all the confidentiality constraints.

Tool for the design of fragmented schemas

Data outsourcing is emerging today as a successful paradigm allowing individuals and organizations to exploit external services for storing and managing huge data collections. A common practice for protecting outsourced data consists in encrypting the data. As described in the previous section, there are situations where encryption may be an overdue and may not always possible. These proposals guarantee the protection of outsourced data by combining encryption with fragmentation. The idea is that sensitive associations among data are protected by splitting them into two or more fragments while information that is sensitive per se is encrypted.

The tool developed in WP2.3 and described in H2.3.4 has the goal of providing a support to the data owner in the fragmentation process. Given a set of confidentiality constraints representing the sensitive data or associations that need to be protected, the tool computes a fragmentation that satisfies the confidentiality constraints. The frag-

mentation algorithm departs from the use of encryption and exploits the availability of a (limited) trusted storage at the data owner side. The tool works under the assumption that the owner, while outsourcing the major portion of the data at one or more external servers, is then willing to locally store a limited amount of data. The owner-side storage, being under the owner control, is assumed to be maintained in a trusted environment. The tool produces two fragments: one stored at the data owner side and the the one stored at the external server. This permits to break sensitive associations by exploiting fragmentation only without using encryption. Since the trusted storage available at the data owner is limited in size and expensive (also because the data owner would be involved in the evaluation of users' queries), the tool has been designed to compute a fragmentation that minimizes the owner's workload. The tool supports different metrics for the evaluation of the quality of a fragmentation. In particular, it is offered the possibility to minimize: (1) the number of attributes stored at the owner; (2) the size of the fragment stored at the owner; (3) the number of queries partially evaluated by the owner; (4) the number of conditions evaluated by the owner.

The tool is composed of two applications: the first implements a greedy fragmentation algorithm (developed in C++), while the second realizes its Graphical User Interface (developed in Java).

3.2.3 Efficient organization and access to privacy-preserving data collections (Task 2.3.3)

The goal of this task is the investigation of techniques able to support efficient access to large collections of private information. The work in the second year has considered a scenario with outsourced data and privacy constraints as defined in Task 2.3.2 in the first year of the project. The goal has been the identification of a fragmentation of data able to optimize query performance.

Optimization of fragmented schemas based on query profiles

The scenario considers a few cases. In a first case, a medical organization must manage a collection of data recording the medical histories of a community of patients. Researchers can then access these data and effectively and efficiently discover behavioral and social patterns that exhibit correlation with specific pathologies, with a direct positive impact on medical research. The downside is that a compromise of the server can disclose patients' information and violate their privacy. In another case, the owner of an e-commerce Web site must store the complete description of the financial data about transactions executed on the site. The Web site offers a wider choice and lower prices than a brick-and-mortar store, producing an immediate benefit to consumers and a considerable positive economic impact. The downside is that a compromise of the Web server may bring customers' data into the black market, where they can be used in fraudulent transactions. The two scenarios demonstrate that, while information and communication technology can provide important benefits, they inevitably introduce risks of exposing private information to improper disclosure.

The crucial observation behind the approach considered in the second year is that users of the system may normally need to access data combining sensitive and non

sensitive information in a way that does not introduce privacy risks. For instance, medical researchers may typically need to access generic and not-identifying patient data when performing their research. The owner of the Web site mostly accesses the financial data about the transactions managed by the Web site with no need to reference the personal data of the customer. On the other hand, medical researchers may sometimes need to evaluate parameters that may lead to the specific identity of the patient, and the Web site owner may need to retrieve the complete credit card data when a dispute arises. In addition, regulations are forcing requirements on the management of personal information that often explicitly demand the use of encryption for the protection of sensitive data.

A promising approach to protect sensitive data or sensitive associations among data stored at external parties is represented by the combined use of fragmentation and encryption, as considered in Task 2.3.2 in the first year of the project. Fragmentation and encryption provide protection of data in storage, or when disseminated, ensuring no sensitive information is disclosed neither directly (i.e., present in the database) nor indirectly (i.e., derivable from other information in the database). With this design, the data can be outsourced and stored on an untrusted server, typically obtaining lower costs, greater availability, and more efficient distributed access. This scenario resembles the "database-as-a-service" (DAS) paradigm [CDD+05, HIML02] and indeed the techniques explored in this task can be considered an adaptation of this paradigm to a context where only part of the information stored into the database is confidential and where the confidentiality of associations among values is protected by storing them in separate fragments. The advantage of having only part of the data encrypted is that all the queries that do not require to reconstruct the confidential information will be managed more efficiently and securely. This approach represents for the real-world database and security administrators a more interesting solution compared to the canonical full-encryption DAS scenario, where the use of the secret key for each access creates a significant vulnerability.

The combined use of fragmentation and encryption to protect confidentiality has been initially proposed in [ABG+05]. However, the proposal in [ABG+05] assumes information to be stored on two separate servers and protection relies on the hypothesis that the servers cannot communicate. This assumption is clearly too strong in any practical situation. The work in this task, presented in [CDF+09b], produced an approach for the design of a fragmentation that looks carefully at the performance issues and takes into account the profile of the query load on the server. A heuristic algorithm has been designed for producing, given a set of confidentiality constraints to be satisfied, a fragmentation design exhibiting good performance. The experimental results support the quality of the solutions produced by the heuristic.

3.3 Future research

3.3.1 Privacy assessment and privacy metrics (Task 2.3.1)

Considering the lines of investigation presented earlier, the plan for Task 2.3.1 is to continue the investigation on the definition of privacy metrics for business applications, and on the definition of novel protection techniques for large data collections.

With respect to privacy metrics for business applications, the next step will be trying

to adapt privacy metrics to the scenarios presented before. The work in Task 2.3.1 will also investigate, with consortium partners, how novel metrics may be defined to address the above mentioned issues.

The work on protection techniques will consider how the fragmentation of a relational schema and the partial reconstruction of the relationships among tuples in the fragment can satisfy the privacy requirements, at the same time offering a significant level of utility in the protected data. The definition of the technique and the investigation of its properties, both in terms of design, protection, and utility, presents several interesting aspects, that deserve to be investigated.

3.3.2 Techniques for enforcing data privacy (Task 2.3.2)

All the research topics studied within Task 2.3.2 in the first two years of the project will be extended.

The investigation on the management of privacy constraints in relational databases will consider richer definitions of constraints, extending the power of the logical language. An interesting potential lies in the definition of an approach that integrates in a single language the representation of privacy constraints and the visibility requirements set by the application. The goal would then be the definition of an approach able to build, in an automatic or semi-automatic way, a system configuration able to support both components of the model, possibly maximizing some performance metric on the cost required for the processing of a predefined query load on the system.

The research on techniques for privacy-conscious business applications will identify which methods are better adapted to the business scenarios considered in the project. As in the case of privacy metrics, particular attention will be devoted to performance.

3.3.3 Efficient organization and access to privacy-preserving data collections (Task 2.3.3)

The work in Task 2.3.3 will first aim to define novel techniques able to support efficient access to protected data, without violating the privacy of access. The task will carefully consider the scenarios and the techniques developed in the other tasks and work packages, aiming to a solution able to work together with the other tools designed in the project.

The design of the technique will be described in papers that will be submitted to scientific conferences and journals. An additional prototype is also planned, to demonstrate the behavior of the technique. An interesting potential can be identified in the design of an indexing structure able to realize efficient access to encrypted data, at the same time hiding the profile of accesses to the data.

Chapter 4

Access control for the protection of user-generated data (WP 2.4)

4.1 Introduction

The widespread access to information and communication channels provided by modern technology has provided significant benefits in our life, since we can enjoy a wide variety of electronic services, which however often require the release of our personal information. The vast amount of personal information thus available has introduced many concerns about the privacy of the users. This work package focuses on the technological aspects of privacy within today's global network infrastructure, where users interact with remote information sources for retrieving data or for using on-line services. The general objective of the work package is the definition of novel models and techniques for enforcing access control restrictions on user-generated data collected and disseminated by external servers. The goal is then to enhance the user awareness and empowerment, granting users the ability to participate in (and be aware of) the management and dissemination of their data and resources. To achieve this goal, the research work in this work package is continued along the research directions identified in the first year of the project. In particular, the main objectives that have been pursued can be briefly summarized as follows.

- *Task 2.4.1.* Definition of novel solutions for limiting the disclosure of personal user information, including information about the context of users (e.g., location information).

- *Task 2.4.2.* Definition of efficient techniques for enforcing access control in a scenario where the personal information of the users is collected and shared by an external (honest-but-curious) server.

- *Task 2.4.3.* Definition of novel techniques for the management of personal information stored in trusted services, for example, in cloud computing.

In the second year of the project most attention has been devoted on Task 2.4.1 and Task 2.4.2. The work in Task 2.4.3 has been low due to an internal project reorganization that resulted in a reduction of the resources dedicated to it. The work in this task has terminated this year.

The research activity performed during the second year of the project resulted in eight publications: two appeared in international journals; four appeared in international conferences and workshops; two are book chapters. In the remainder of this chapter, we first describe the advancement status of the research work done in the second year of the project and then outline future work.

4.2 Research results

The results obtained in the second year of the project are well aligned with the overall goal of the work package. In addition to the research contributions that will be described in the following, the work package has also produced the internal heartbeat H2.4.4 that presents a tool demonstrating the techniques for the realization of access control policies with encryption. For each task composing the work package, we now provide a brief description of the main research contributions.

4.2.1 Dissemination control and secondary use restrictions (Task 2.4.1)

The work in this task addressed the problem of defining a powerful solution for regulating the access and dissemination of different types of personal information, including location information. To protect the privacy of the users to whom the information collected and disseminated refers, it is fundamental to define powerful access control models, policies, and languages supporting privacy requirements [ABC09] and to define novel techniques for protecting the secondary use of context information [ACDS09]. These problems have been studied from different perspectives, resulting in the following contributions.

XML secure views using semantic access control

Internet technologies appear to facilitate the information flow between entities. There are however risks connected to the misuse of data belonging either to organizations or to individuals. In fact, collaborations between two or more subjects usually involve disclosure of information, which, if it is not controlled, might end in revealing strategic and valuable company knowledge or sensitive data about citizens, workers or clients. With data protection and privacy acts, such as the EU directive 95/46/EC [Eur95], governments have been trying to regulate ways in which data about citizens can be collected, used and transferred, allowing the individuals to express preferences about the usage of their own data, using privacy policy.

Moving towards a system that is able to selectively disclose information, according to regulatory requirements, we focused on a solution able to alter out data from XML documents before they are sent or given away, according to policies that describe the rules in terms of whether a concept may be released or not. The altering is based on policies written using the eXtensible Access Control Mark-up Language (XACML),

which has been modified to interact with an ontology knowledge base. Thanks to this approach, systems can regulate the disclosure of parts of the document not only by schema constraints, but also by the semantics of their contents. This proposal has been presented in a publication [SRR10].

Context information and location privacy

Context information should be used by the policy infrastructure to allow environment factors to influence how and when policy is enforced. As far as policy enforcement is concerned, context contains information enabling verification of policy conditions and, therefore, it should be made available to any authorized service/application at any time and in a standard format. Still unauthorized information leaks should be prevented, also to avoid loss of privacy, for example, on the user's whereabouts.

The first contribution toward the exploitation of context information is related to the use of this kind of information for detecting and monitoring the presence of users [DAP09, DCP09]. The possibility of correctly managing this information can be exploited to define powerful access control policies, where the information about the presence of users can influence the actions that can be performed.

The second contribution is focused on a crucial type of context information, that is, the *location information* that makes it possible to develop context-sensitive services, where access to resources provided/managed by a server is limited depending on a user's context. For instance, a *location-based service* (LBS) can require a user to be at a particular location to let the user use or access a resource or learn her friends' location. Customer-oriented applications, social networks, and monitoring services can then be greatly enriched with data reporting where people are, how they are moving, or whether they are close to specific locations. Several commercial and enterprise-oriented LBSs are already available and have gained popularity (e.g., [BD03, DB03, Loo08]), driven by the relevant enhancements achieved in the field of sensing technologies. Location techniques permit to gather location information with good precision and reliability at costs that most people (e.g., the cost of current mobile devices like cellular phones) and companies (e.g., the cost of integrating location techniques in current telecommunication systems) can economically sustain.

In this context, the privacy of the users, which is already the center of many concerns for the risks posed by current online services (e.g., [BD03, Pri06]), can be threatened by LBSs. The publicity gained by recent security incidents that have targeted the privacy of users has revealed faulty data management practices and unauthorized trading of personal information (including ID thefts and unauthorized profiling). For instance, legal cases have been reported, where rental companies used the GPS technology to track their cars and charge users for agreement infringements [Chi01], or where an organization used a location service to track its own employees [Lee04]. In addition, research on privacy issues has gained a relevant boost since providers of online and mobile services have often largely exceeded in collecting personal information in the name of service provision.

In such a worrisome scenario, the concept of *location privacy* can be defined as *the right of individuals to decide how, when, and for which purposes their location information can be released to other parties.* The improper exposure of location information could result in severe consequences that make users the target of fraudulent attacks [DK06].

In the second year of PrimeLife we then worked on a novel solution, presented in a publication [ACDS10], aimed at preserving the location privacy of the users by perturbing location information measured by sensing technologies. We focused on the development of novel techniques for protecting a single sample of location information. For the sake of concreteness, we considered locations gathered by means of cellular phones as our reference, even if our solution is not bound to a specific location technique. One important characteristic of cellular phones is their large availability and the possibility to be used as a source of location information both indoor and outdoor (on the contrary, GPS is operating mainly outdoor).

The location information of users, as each physical measurement, is always affected by an intrinsic measurement error introduced by sensing technologies. A direct consequence of such a lack of precision is that the location position of a user cannot be expressed as a geographical point, which would imply to suppose that sensing technologies can return exact information. We then assume that positions of users are always represented as *planar circular areas*. This assumption satisfies the general requirement of considering convex areas to easily compute integrals over them. Also, circular areas approximate well the actual shape resulting from many location techniques (e.g., location gathering based on cellular phones). A *location measurement* returned by a sensing technology is therefore characterized by the coordinates of its center and by its radius. We also assume that: *(1)* the real user position is within the returned location measurement; and *(2)* the real user position could be randomly located everywhere inside the location measurement with uniform probability.

The key aspects of our perturbation process, called *obfuscation*, are to allow users to express their privacy preferences in a simple and intuitive way, and to enforce the privacy preferences through a set of *obfuscation techniques* robust against a relevant class of de-obfuscation attacks. To this end, we introduce the concept of *relevance* as a metric of both location information accuracy and privacy that abstracts from the physical attributes of the sensing technology as well as from the actual technique employed to obfuscate a location. In this way, while users have just to select a relevance value, the robustness of the solution is guaranteed by randomly selecting one of the techniques to produce the obfuscated location. With respect to the obfuscation techniques, we have defined the following three basic techniques (see Figure 9).

- *Enlarge* (`E`). Given a location measurement with radius r, it is obfuscated by enlarging the radius (i.e., from r to r' with $r' > r$). Obfuscating a location measurement by increasing its radius logically corresponds to generalization techniques employed in data privacy solutions (e.g., [CDFS07]). Intuitively, such an obfuscation has the effect of decreasing the probability that the real user position falls within the obfuscated location measurement.

- *Shift* (`S`). Given a location measurement with center (x, y), it is obfuscated by shifting the center (i.e., from (x, y) to $(x + d\sin\theta, y + d\cos\theta)$, where d is the distance between the centers of the original and obfuscated measurement and θ is a rotation angle). Such an obfuscation has the probabilistic effect of decreasing both the probability that the real user position is in the neighborhood of a point of the obfuscated area and the probability that the real user position falls within the obfuscated area.

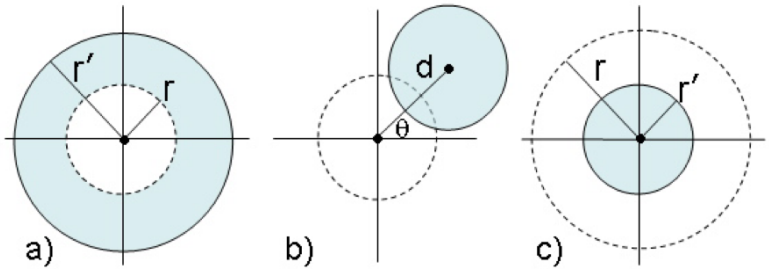

Figure 9: Enlarging (a), shifting (b), and reducing (c)

- *Reduce* (`R`). Given a location measurement with radius r, it is obfuscated by reducing the radius (i.e., from r to r' with $r' < r$). While this obfuscation effect might appear counter-intuitive at first sight, it has a precise probabilistic explanation: the probability that the real user position falls within the obfuscated area is reduced.

The physical transformations resulting from these three basic obfuscation operators can be composed by executing them in sequence. Although in theory it is possible to combine operators `E`, `R`, and `S` an arbitrary number of times, the combination of more than two operators is never necessary.

The robustness of the basic and composed operators has been experimentally evaluated by simulating the adversary behavior under different assumptions: when the adversary is not aware of any contextual information and is not able to infer the obfuscation family applied; when the adversary knows enough contextual information to infer the obfuscation family applied. During the tests we measured the robustness of our operators, compared one with the others, and with the trivial solution based on just an enlargement of the location measurement. Quantitative evaluations and comparisons are produced based on both the successful de-obfuscation rate achieved by the adversary and the relevance gain or loss the adversary obtains as a result of a de-obfuscation attempt. Analyzing both aspects (i.e., the success rate and the amount of gain/loss) is relevant because in a real scenario the adversary is assumed to behave strategically by, implicitly or explicitly, maximizing them. Our experiments demonstrate that the obfuscation techniques are robust against attackers aiming at reversing the protection granted by obfuscation.

4.2.2 Access control to confidential data stored at external services (Task 2.4.2)

The work in this task has been focused on the definition of techniques and models for efficiently enforcing selective access to sensitive information when it is outsourced to external servers with the goal of minimizing the overhead it causes to the users. Data outsourcing has become increasingly popular in recent years [HIML02]. In fact, users are more and more interested in sharing and disseminating their personal information using

the services provided by external parties (e.g., Web sites such as Facebook or MySpace have millions of users using their services). Companies are also interested in exploiting external services for managing their (potentially sensitive) data since the design, realization, and management of a secure system able to grant the confidentiality of sensitive information might be very expensive. Data outsourcing presents important advantages: management costs are reduced and higher availability and more effective disaster protection than in-house operations are provided. On the other hand, data outsourcing opens the door to possible violations to the data management requirements and introduces therefore new issues to be addressed. Being stored externally, data are not under the control of their owners anymore, their confidentiality and integrity can therefore be at risk. This aspect is also recognized by recent regulations (e.g., California Senate Bill SB 1386 and the Italian Personal data protection code, Legislative Decree no. 196 of 30 June 2003) that explicitly require specific categories of sensitive information to be either *encrypted* or *kept separate* from other personally identifiable information to ensure data confidentiality. In many cases, the server itself might not be allowed to read the actual content of the data outsourced to it for storage and management. In this case, the *honest-but-curious* server should provide effective service while operating on data that should result not intelligible to it. While trustworthy with respect to their services in making outsourced information available, these external servers are however trusted neither to access the content nor to fully enforce access control policy and privacy protection requirements. It is therefore of primary importance to provide means of protecting the confidentiality of the information remotely stored, without necessarily requiring trust in the subject managing the information, while guaranteeing its availability to legitimate users.

Since the enforcement of the authorization policy cannot be delegated to the remote server, the data owner has to be involved in the access control enforcement. To avoid the owner's involvement in managing access and enforcing authorizations, recently *selective encryption* techniques have been proposed [DDF$^+$06, DFJ$^+$07]. Selective encryption means that the *encryption policy* (i.e., which data are encrypted with which key) is dictated by the authorizations to be enforced on the data. The basic idea is to use different keys for encrypting data and to release to each user the set of keys necessary to decrypt all and only the resources the user is authorized to access. To avoid users having to store and manage a huge number of (secret) keys, these proposals are based on a *key derivation method* that allows the derivation of a key starting from another key and some public information [AT83, AFB05, ADFM06, DFM04, MTMS85, San88]. The derivation relationship between keys can be represented through a *user graph*, where there is a vertex v for each possible set of users and a key associated with it, and there is an edge (v_i,v_j) for all pairs of vertices such that the set of users represented by v_i is a subset of the set of users represented by v_j. The authorization policy can be enforced: *i)* by encrypting each resource with the key of the vertex corresponding to its access control list, and *ii)* by assigning to each user the key associated with the vertex representing the user in the graph. Since edges represent the possible key derivations, each user u, starting from her own key, can directly compute the keys of all vertices v such that u belongs to the group represented by v. It is easy to see that this approach to design the encryption policy *correctly enforces* the authorization policy defined by the owner.

Along this line of research, an interesting problem that has been investigated in this

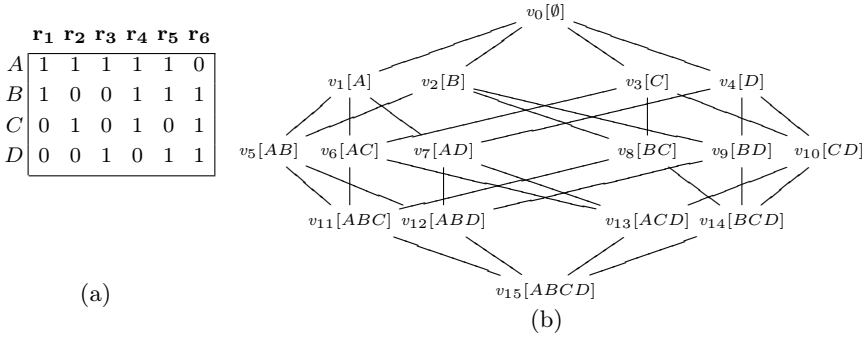

Figure 10: An example of access matrix (a) and of user graph over users $\{A,B,C,D\}$ (b)

task is the minimization of the number of keys to be maintained by the system and distributed to users. Indeed, although the solution based on a user hierarchy is conceptually simple and potentially easy to implement, it defines significantly more keys than actually needed. Furthermore, a crucial aspect for the success of a solution supporting selective encryption is the efficiency of the key management and distribution activities required. Since key derivation methods working on trees are in general more convenient and simpler than those working on directed acyclic graphs and require a lower amount of publicly available information, the solution to be designed should be based on a transformation of the user graph into an equivalent *user tree*, able to enforce the same access control policy.

The problem of transforming a user graph is a *user tree* and of minimizing of the number of keys to be maintained by the system and distributed to users has been presented in two publications [BCD+09, BCD+10]. The idea is that the user tree is defined on a subset of the vertices in the user graph, thus reducing the number of keys in the system. In particular, we define a user tree as a tree, subgraph of the user graph, rooted at the vertex representing the empty group of users and spanning all the vertices, called *material vertices*, representing sets of users corresponding to the access control lists of the resources to be protected (i.e., vertices whose keys are used for encrypting resources). To guarantee the correct enforcement of the authorization policy, each user u is communicated a *key ring*, including the keys associated with any vertex v such that u belongs to the group represented by v, but not to the group of users represented by its parent in the tree. To better illustrate the idea behind our approach, consider the authorization policy illustrated in Figure 10(a). Here, the authorization policy defined by the data owner is represented through an *access matrix* $\mathcal{A}$, with a row for each user u in the system (users A, B, C, and D in the example), a column for each resource r to be protected (resources $r_1 \ldots r_6$, in the example), and $\mathcal{A}[u, r]$ equal to 1 (0, resp.) if u has (does not have, resp.) authorization to access r. Figure 10(b) illustrates an example of user graph where for each vertex v_i, the users in the square brackets correspond to the set of users represented by the vertex. It is easy to see that the user graph correctly enforces the

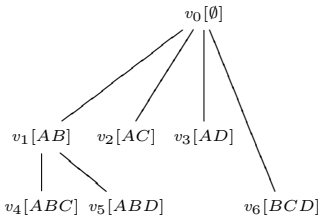

(a)

user	**key_ring**
A	$v_1.key$, $v_2.key$, $v_3.key$
B	$v_1.key$, $v_6.key$
C	$v_2.key$, $v_4.key$, $v_6.key$
D	$v_3.key$, $v_5.key$, $v_6.key$

(b)

Figure 11: A user tree (a) and the corresponding key rings (b)

given authorization policy. For instance, user A can use her key (i.e., the key associated with vertex v_1) to derive the keys associated with vertices $v_5, \ldots, v_7, v_{11}, \ldots, v_{13}$, and v_{15}. According to the authorization policy in Figure 10(a), user A is authorized to access resources $r_1, \ldots, r_5$. Since these resources are encrypted with the keys associated with the vertices representing their access control lists (i.e., vertices v_5, v_6, v_{11}, and v_{12}), they are all accessible to A. Figure 11(a) illustrates an example of user tree correctly enforcing the authorization policy in Figure 10(a), and Figure 11(b) reports the corresponding user key rings (notation $v.key$ denotes the key associated with vertex v). The user tree clearly requires the management of a number of keys that is less than the number of keys needed with a user graph (i.e., a key for each vertex of the user graph).

Given a set of users and an authorization policy, more user trees may exist. We are however interested in determining a user tree correctly enforcing a given authorization policy and minimizing the number of keys in users' key rings. Since the problem of computing a minimum user tree is NP-hard, we propose two alternative approaches. First, we formulate the minimization problem as an *Integer Linear Programming* (ILP) problem, which can be solved adopting known algorithms and tools. Second, we introduce *three families of heuristics*, which are based on the computation of a minimum spanning tree induced by material vertices over the user graph. All the proposed heuristics are based on the observation that the weight of a user tree spanning material vertices can be reduced with the addition of *non-material* vertices, representing sets of users resulting from the intersection of the groups of users associated with at least two vertices already in the tree. Consider, as an example, two vertices v_i and v_j in the tree. The possibly insertion of a new vertex v_k (representing the group users belonging to both the groups represented by v_i and v_j) as a parent of v_i and v_j can reduce the weight of the tree. Indeed, the key ring of the users in the group represented by v_k should only include the key associated with v_k, instead of both the keys associated with v_i and v_j. The three families of heuristics differ in the strategy adopted to individuate non-material vertices for reducing the cost of the minimum spanning tree.

We experimentally evaluated the three families of heuristics, to provide the system designer with a valid set of tools she can use for the selection of the strategy that provides the best trade-off between the quality of the solution returned by the selected heuristic and the amount of time invested in obtaining such a result. These results have been compared with both the solution computed by the algorithm proposed in [DDF$^+$06] and

the one computed through the integer linear programming formulation of the problem. The experimental results obtained by the implementation of the heuristics prove their effectiveness and their efficiency with respect to previous solutions. Also, the experiments performed confirm that the time necessary to solve the ILP problem is always higher than the time necessary to execute one of the three families of heuristics, while the weight of the user tree is often very close to the weight of the trees computed by the heuristics.

4.2.3 User-managed access control to personal data stored in trusted services (Task 2.4.3)

The work in this task started to look at the management of user personal information stored in "trusted" services (e.g., in cloud computing). Trusted services are services that will behave as expected in terms of enforcing data handling, access control policies, and fulfilling obligations. The work in this task has been, however, low in the second year of the project since the main focus has been set on the management of obligations, which mainly fall in WP5.2.

4.3 Future research

The research results obtained in the second year of PrimeLife are important steps towards the main goal of Work Package 2.4. The work package will contribute with other research results in the remaining year. For the first two tasks of the work package, we now outline future work. Note that Task 2.4.3 is terminated and no further work is therefore planned for it.

4.3.1 Dissemination control and secondary use restrictions (Task 2.4.1)

Although users provide personal information for use in one specific context, they often have no idea on how such a personal information may be used subsequently. In other words, users do not always realize that the information they disclose for one purpose (e.g., name, date of birth, and address within an on-line transaction) may be subsequently used for different purposes. Our goal for the last year of the project will be then the development of novel solutions effectively granting more control to the users on the specification and enforcement of access restrictions on their personal data, and considering the privacy requirements of the users when their personal information is collected for a specific purpose (e.g., for accessing a given service).

4.3.2 Access control to confidential data stored at external services (Task 2.4.2)

Within the reference scenario described in Section 4.2.2, there are different open issues that we plan to address in the future research.

A first important problem is that our access control system, based on selective encryption, manages access control enforcement under the assumption that access operations are *read only*. This assumption, even if adequate in a data dissemination scenario, is

not sufficient for the management of data subject to dynamic updates by different parties, which may not coincide with the data owner. In the *multi-owner* scenario, each owner is authorized to modify the portion of data she owns, while she can only read a larger subset of the outsourced resources, possibly owned by another party. We plan then to extend our access control model, relaxing the assumption that accesses are read-only, and proposing a system able to efficiently manage also the multi-owner scenario. Current works on integrity in the data outsourcing scenario, while guarantee that write operations are performed by authorized users only, are not suited to the multi-owner scenario, since they do not allow different users to have write privileges on different subsets of resources (i.e., each user has write privileges on the resources she own). Therefore, we plan to define a model based on a user graph/tree, obtained by merging the graphs representing the policies defined by the different data owners in the system. We will evaluate the robustness of the approach against attacks coming both from internal users and from external attackers.

Another crucial point is the definition of a formal model for representing and transforming an authorization policy defined by the data owner through an equivalent encryption policy. This model will be guided by the principles of releasing at most one key to each user, and encrypting each resource at most once. Like for our previous proposals, we plan to exploit a hierarchical organization of keys allowing the derivation of keys from other keys and public tokens. Our goal is then to minimize the number of tokens to be generated and maintained. We will also address the problem of enforcing updates to the authorization policy while limiting the cost in terms of bandwidth and computational power (providing a two layer approach that avoids the need for the owner to download the affected resources, decrypt and re-encrypt them, and reload their new versions). It is important to note our goal is to develop a solution that should be independent from any specific data model and that should not rely on any specific authorization language.

Chapter 5

Conclusions

During the second year of PrimeLife, the work packages of Activity 2 have made, and continue to make, significant contributions to the development of privacy-aware techniques for ensuring privacy and trust in the electronic society. WP2.1 has improved the state-of-the-art in different cryptographic areas related to anonymity and to the protection of the users privacy in general. WP2.2 has worked towards the analysis of: transparency support tools; privacy measurements for evaluating the degree of privacy provided to the users; and privacy aspects regarding collaborative groups and reputation systems. WP2.3 has developed new solutions for measuring and protecting privacy of large data collections, paying also attention to the identification of efficient configurations. WP2.4 has considered the issue of access control to location information and investigated dynamic access control mechanisms that can be driven by users for providing access to their data.

The reported research results have been presented at and published in the proceedings of different leading international conferences, including CRYPTO, ACM CCS, EUROCRYPT, ESORICS, IEEE ICDCS.

The research results presented in this report represent a significant step towards the realization of the overall goal of PrimeLife. In the third year, the research will continue towards the development of novel and innovative solutions for ensuring privacy of the users in the electronic society. Open issues that will be investigated include: the development of solutions for an efficient revocation of anonymous credentials; the development of mechanisms for supporting anonymous credentials in light-weight devices; the design of novel protocols for private service access; the design of privacy aware third-party services; the advancement of trusted wallet technologies; the evolution of log-based systems; the design of entropy-based privacy metrics; the evolution of tools for privacy-aware group collaboration; the enhancement of user awareness; the definition of privacy techniques for large data collections; the investigation of richer definitions of privacy constraints in relational databases; the design of novel indexing techniques for encrypted data; the definition of a formal model for transforming an authorization policy into an equivalent encryption policy as well as the development of a solution supporting policy changes in a scenario where the information is stored at an external server; and, the

definition of an access control solution for a multi-user environment.

Chapter 6

Abstracts of research papers

6.1 Cryptographic mechanisms (WP 2.1)

1. J. Camenisch, M. Kohlweiss, and C. Soriente, "An Accumulator Based on Bilinear Maps and Efficient Revocation for Anonymous Credentials," in *Public Key Cryptography, 12th International Workshop on Practice and Theory in Public Key Cryptosystems (PKC 2009)* [CKS09].
   Abstract. The success of electronic authentication systems, be it e- ID card systems or Internet authentication systems such as CardSpace, highly depends on the provided level of user-privacy. Thereby, an important requirement is an efficient means for revocation of the authentication credentials. In this paper we consider the problem of revocation for certificate-based privacy-protecting authentication systems. To date, the most efficient solutions for revocation for such systems are based on cryptographic accumulators. Here, an accumulate of all currently valid certificates is published regularly and each user holds a witness enabling her to prove the validity of her (anonymous) credential while retaining anonymity. Unfortunately, the users' witnesses must be updated at least each time a credential is revoked. For the know solutions, these updates are computationally very expensive for users and/or certificate issuers which is very problematic as revocation is a frequent event as practice shows.

   In this paper, we propose a new dynamic accumulator scheme based on bilinear maps and show how to apply it to the problem of revocation of anonymous credentials. In the resulting scheme, proving a credential's validity and updating witnesses both come at (virtually) no cost for credential owners and verifiers. In particular, updating a witness requires the issuer to do only one multiplication per addition or revocation of a credential and can also be delegated to untrusted entities from which a user could just retrieve the updated witness. We believe that thereby we provide the first authentication system offering privacy protection suitable for implementation with electronic tokens such as eID cards or drivers' licenses.

2. J. Camenisch, M. Kohlweiss, A. Rial, and C. Sheedy, "Blind and Anonymous Identity-Based Encryption and Authorised Private Searches on Public Key Encrypted Data," in *Public Key Cryptography, 12th International Workshop on Practice and Theory in Public Key Cryptosystems (PKC 2009)* [CKRS09].
   Abstract. Searchable encryption schemes provide an important mechanism to cryptographically protect data while keeping it available to be searched and accessed. In a common approach for their construction, the encrypting entity chooses one or several keywords that describe the content of each encrypted record of data. To perform a search, a user obtains a trapdoor for a keyword of her interest and uses this trapdoor to find all the data described by this keyword.

   We present a searchable encryption scheme that allows users to privately search by keywords on encrypted data in a public key setting and decrypt the search results. To this end, we define and implement two primitives: public key encryption with oblivious keyword search (PEOKS) and committed blind anonymous identity-based encryption (IBE). PEOKS is an extension of public key encryption with keyword search (PEKS) in which users can obtain trapdoors from the secret key holder without revealing the keywords. Furthermore, we define committed blind trapdoor extraction, which facilitates the definition of authorisation policies to describe which trapdoor a particular user can request. We construct a PEOKS scheme by using our other primitive, which we believe to be the first blind and anonymous IBE scheme. We apply our PEOKS scheme to build a public key encrypted database that permits authorised private searches, i.e., neither the keywords nor the search results are revealed.

3. Jan Camenisch, Maria Dubovitskaya, Gregory Neven, "Oblivious transfer with access control," in *ACM Conference on Computer and Communications Security 2009* [CDN09]
   Abstract. We present a protocol for anonymous access to a database where the different records have different access control permissions. These permissions could be attributes, roles, or rights that the user needs to have in order to access the record. Our protocol offers maximal security guarantees for both the database and the user, namely (1) only authorized users can access the record; (2) the database provider does not learn which record the user accesses; and (3) the database provider does not learn which attributes or roles the user has when she accesses the database.

   We prove our protocol secure in the standard model (i.e., without random oracles) under the bilinear Diffie-Hellman exponent and the strong Diffie-Hellman assumptions.

4. Jan Camenisch, Maria Dubovitskaya, Gregory Neven, "Unlinkable Priced Oblivious Transfer with Rechargeable Wallets," in *Finanical Cryptography 2010* [CDN10]
   Abstract. We present the first truly unlinkable priced oblivious transfer protocol. Our protocol allows customers to buy database records while remaining fully anonymous, i.e., (1) the database does not learn who purchases a record, and cannot link purchases by the same customer; (2) the database does not learn which record is being purchased, nor the price of the record that is being purchased; (3) the customer can only obtain a single record per purchase, and cannot spend more

than his account balance; (4) the database does not learn the customer's remaining balance. In our protocol customers keep track of their own balances, rather than leaving this to the database as done in previous protocols. Our priced oblivious transfer protocol is also the first to allow customers to (anonymously) recharge their balances. Finally, we prove our protocol secure in the standard model (i.e., without random oracles).

5. Patrik Bichsel, Jan Camenisch, Thomas Groß, Victor Shoup, "Anonymous credentials on a standard java card. " in *ACM Conference on Computer and Communications Security 2009* [BCGS09]
   Abstract. Secure identity tokens such as Electronic Identity (eID) cards are emerging everywhere. At the same time user- centric identity management gains acceptance. Anonymous credential schemes are the optimal realization of user-centricity. However, on inexpensive hardware platforms, typically used for eID cards, these schemes could not be made to meet the necessary requirements such as future- proof key lengths and transaction times on the order of 10 seconds. The reasons for this is the need for the hardware platform to be standardized and certified. Therefore an implementation is only possible as a Java Card applet. This results in severe restrictions: little memory (transient and persistent), an 8-bit CPU, and access to hardware acceleration for cryptographic operations only by defined interfaces such as RSA encryption operations. Still, we present the first practical implementation of an anonymous credential system on a Java Card 2.2.1. We achieve transaction times that are orders of magnitudes faster than those of any prior attempt, while raising the bar in terms of key length and trust model. Our system is the first one to act completely autonomously on card and to maintain its properties in the face of an untrusted terminal. In addition, we provide a formal system specification and share our solution strategies and experiences gained and with the Java Card.

6. Patrik Bichsel, Samuel Müller,Franz-Stefan Preiß, Dieter Sommer, Mario Verdicchio, "Security and Trust through Electronic Social Network-Based Interactions" in *CES 2009* [BMP$^+$09] **Abstract.** The success of a Public Key Infrastructure such as the Web of Trust (WoT) heavily depends on its ability to ensure that public keys are used by their legitimate owners, thereby avoiding malicious impersonations. To guarantee this property, the WoT requires users to physically gather, check each other's credentials (e.g., ID cards), to sign the trusted keys, and to subsequently monitor their validity over time. This trust establishment and management procedure is rather cumbersome and, as we believe, the main reason for the limited adoption of the WoT. To overcome this problem, we propose a solution that leverages the intrinsic properties of Electronic Social Networks (ESN) to establish and manage trust in the WoT. In particular, we exploit dynamically changing profile and contact information, as well as interactions among users of ESNs to gain and maintain trust in the legitimacy of key ownerships without the disadvantages of the traditional WoT approach. We see our proposal as an effective way to make security and trust solutions available to a broad audience of non-technical users.

7. Mira Belenkiy, Jan Camenisch, Melissa Chase, Markulf Kohlweiss, Anna Lysyan-

skaya, Hovav Shacham, "Randomizable Proofs and Delegatable Anonymous Credentials" in *CRYPTO 2009* [BCC+09]

Abstract. We construct an efficient delegatable anonymous credentials system. Users can anonymously and unlinkably obtain credentials from any authority, delegate their credentials to other users, and prove possession of a credential L levels away from a given authority. The size of the proof (and time to compute it) is O(Lk), where k is the security parameter. The only other construction of delegatable anonymous credentials (Chase and Lysyanskaya, Crypto 2006) relies on general non-interactive proofs for NP-complete languages of size kï¿½(2L). We revise the entire approach to constructing anonymous credentials and identify randomizable zero-knowledge proof of knowledge systems as the key building block. We formally define the notion of randomizable non-interactive zero-knowledge proofs, and give the first instance of controlled rerandomization of non-interactive zero-knowledge proofs by a third-party. Our construction uses Groth-Sahai proofs (Eurocrypt 2008).

8. Jan Camenisch, Nishanth Chandran, Victor Shoup, "A Public Key Encryption Scheme Secure against Key Dependent Chosen Plaintext and Adaptive Chosen Ciphertext Attacks," in *EUROCRYPT 2009* [CCS09]

   Abstract. Recently, at Crypto 2008, Boneh, Halevi, Hamburg, and Ostrovsky (BHHO) solved the long-standing open problem of circular encryption, by presenting a public key encryption scheme and proving that it is semantically secure against key dependent chosen plaintext attack (KDM-CPA security) under standard assumptions (and without resorting to random oracles). However, they left as an open problem that of designing an encryption scheme that simultaneously provides security against both key dependent chosen plaintext and adaptive chosen ciphertext attack (KDM-CCA2 security). In this paper, we solve this problem. First, we show that by applying the Naor-Yung ï¿½double encryptionï¿½ paradigm, one can combine any KDM-CPA secure scheme with any (ordinary) CCA2 secure scheme, along with an appropriate non-interactive zero-knowledge proof, to obtain a KDM-CCA2 secure scheme. Second, we give a concrete instantiation that makes use the above KDM-CPA secure scheme of BHHO, along with a generalization of the Cramer-Shoup CCA2 secure encryption scheme, and recently developed pairing-based NIZK proof systems. This instantiation increases the complexity of the BHHO scheme by just a small constant factor.

9. Jan Camenisch, Aggelos Kiayias, Moti Yung, "On the Portability of Generalized Schnorr Proofs," in *EUROCRYPT 2009* [CKY09]

   Abstract. The notion of Zero Knowledge Proofs (of knowledge) [ZKP] is central to cryptography; it provides a set of security properties that proved indispensable in concrete protocol design. These properties are defined for any given input and also for any auxiliary verifier private state, as they are aimed at any use of the protocol as a subroutine in a bigger application. Many times, however, moving the theoretical notion to practical designs has been quite problematic. This is due to the fact that the most efficient protocols fail to provide the above ZKP properties for all possible inputs and verifier states. This situation has created various problems to protocol designers who have often either introduced imperfect protocols

with mistakes or with lack of security arguments, or they have been forced to use much less efficient protocols in order to achieve the required properties. In this work we address this issue by introducing the notion of ï¿½protocol portability,ï¿½ a property that identifies input and verifier state distributions under which a protocol becomes a ZKP when called as a subroutine in a sequential execution of a larger application. We then concentrate on the very efficient and heavily employed ï¿½Generalized Schnorr Proofsï¿½ (GSP) and identify the portability of such protocols. We also point to previous protocol weaknesses and errors that have been made in numerous applications throughout the years, due to employment of GSP instances while lacking the notion of portability (primarily in the case of unknown order groups). This demonstrates that cryptographic application designers who care about efficiency need to consider our notion carefully. We provide a compact specification language for GSP protocols that protocol designers can employ. Our specification language is consistent with the ad-hoc notation that is currently widely used and it offers automatic derivation of the proof protocol while dictating its portability (i.e., the proper initial state and inputs) and its security guarantees. Finally, as a second alternative to designers wishing to use GSPs, we present a modification of GSP protocols that is unconditionally portable (i.e., ZKP) and is still quite efficient. Our constructions are the first such protocols proven secure in the standard model (as opposed to the random oracle model).

10. Michel Abdalla, Mihir Bellare, and Gregory Neven, "Robust Encryption," in *TCC 2010* [ABN10]
Abstract. We provide a provable-security treatment of "robust" encryption. Robustness means it is hard to produce a ciphertext that is valid for two different users. Robustness makes explicit a property that has been implicitly assumed in the past. We argue that it is an essential conjunct of anonymous encryption. We show that natural anonymity-preserving ways to achieve it, such as adding recipient identification information before encrypting, fail. We provide transforms that do achieve it, efficiently and provably. We assess the robustness of specific encryption schemes in the literature, providing simple patches for some that lack the property. We present various applications. Our work enables safer and simpler use of encryption.

11. Gregory Neven, Nigel Smart, and Bogdan Warinschi, "Hash function requirements for Schnorr signatures," in *Journal of Mathematical Cryptology* [NSW09]
Abstract. We provide two necessary conditions on hash functions for the Schnorr signature scheme to be secure, assuming compact group representations such as those which occur in elliptic curve groups. We also show, via an argument in the generic group model, that these conditions are sufficient. Our hash function security requirements are variants of the standard notions of preimage and second preimage resistance. One of them is in fact equivalent to the Nostradamus attack by Kelsey and Kohno (Eurocrypt 2006), and, when considering keyed compression functions, both are closely related to the ePre and eSec notions by Rogaway and Shrimpton (FSE 2004). Our results have a number of interesting implications in practice. First, since security does not rely on the hash function being collision resistant, Schnorr signatures can still be securely instantiated with SHA-1/SHA-

256, unlike DSA signatures. Second, we conjecture that our properties require $O(2^n)$ work to solve for a hash function with n-bit output, thereby allowing the use of shorter hashes and saving twenty-five percent in signature size. And third, our analysis does not reveal any significant difference in hardness between forging signatures and computing discrete logarithms, which plays down the importance of the loose reductions in existing random-oracle proofs, and seems to support the use of "normal- size" groups.

12. C. Diaz, E. Kosta, H. Dekeyser, M. Kohlweiss, and G. Nigusse, "Privacy preserving electronic petitions," in *Identity in the Information Society* [DKD+09]. **Abstract.** We present the design of a secure and privacy preserving e-petition system that we have implemented as a proof-of-concept demonstrator. We use the Belgian e-ID card as source of authentication, and then proceed to issue an anonymous credential that is used to sign petitions. Our system ensures that duplicate signatures are detectable, while preserving the anonymity of petition signers. We analyze the privacy and security requirements of our application, present an overview of its architecture, and discuss the applicability of data protection legislation to our system.

6.2 Mechanisms supporting users' privacy and trust (WP 2.2)

1. S. Steinbrecher, S. Groß, Markus Meichau, "Jason: A scalable reputation system for the semantic web," in *Proc. of the 24th IFIP TC-11 International Information Security Conference (SEC 2009)* [SGM09].
   Abstract. The recent development of the Internet, especially the expanding use of social software and dynamic content generation commonly termed as Web 2.0 enables users to find information about almost every possible topic on the Web. On the downside, it becomes more and more difficult to decide which information can be trusted in. In this paper we propose the enhancement of Web 2.0 by a scalable and secure cross-platform reputation system that takes into account a user's social network. Our proposed solution *Jason* is based on standard methods of the semantic web and does not need a central entity. It enables the fast and flexible evaluation of arbitrary content on the World Wide Web. In contrast to many other reputation systems it provides mechanisms to ensure the authenticity of web content, thus, enabling the user to explicitely choose information published by trusted authors.

2. Stefan Schiffner, Sebastian Clauß: "Using Linkability Information to Attack Mix-based Anonymity Services," in *Proc. of the 9th Privacy Enhancing Technologies Symposium (PETS 2009)* [SC09]
   Abstract. There exist well established models for anonymity focusing on traffic analysis, i.e., analysing properties of single messages as, e.g., timing. However there is only little work done that use linkability information, that is information about the probability that two messages have been sent by the same sender.

In this paper we model information about linkability between messages as a weighted graph. We show lower and upper bounds with regards to the usefulness of linkability information for matching messages to senders. In addition to that we present simulation results, showing to which extent a matching of messages to senders is possible by using linkability information with different grades of noise.

3. Stefanie Pötzsch, "Untersuchung des Einflusses von wahrgenommener Privatsphäre und Anonymität auf die Kommunikation in einer Online-Community," in *S. Fischer, E. Maehle, and R. Reischuk (Eds.) Informatik 2009, Im Fokus das Leben, 28 September–02 October 2009, Lübeck, volume 154 of Lecture Notes in Informatics, pp. 2152–2165. Gesellschaft für Informatik, Bonn, 2009.* [Pöt09b].
Abstract. The Web 2.0 provides a variety of social software and collaborative applications which serve as a platform for computer-mediated communication and social interactions. Existing studies show, that people subjectively perceive a high level of privacy when using these applications and thus, are willing to disclose a lot of personal data. Against the backdrop of the cues-filtered-out theory, we argue that we can influence people's perception of privacy and finally their disclosing behaviour by providing privacy-related cues. In this paper we present an emprical study to research the effect of privacy-awareness information on people's perception of privacy during their work with a forum and on their actual behavior. (Full paper available in German)

4. S. Schiffner, S. Clauß, S. Steinbrecher, "Privacy and Liveliness for Reputation Systems," in *Proc. of the 6th European PKI Workshop: Research and Applications (EuroPKI 2009)* [SCS09].
Abstract. Privacy-respecting reputation systems have been constructed based on anonymous payment systems in order to implement raters' anonymity. To the best of our knowledge, all these systems suffer from the problem of having a "final state", i.,e., a system state in which users have no incentive anymore to behave honestly because they reached a maximum reputation or they can no longer be rated. Thus the reputation is in fact no longer lively. We propose a novel approach to address the problem of liveliness by the employment of negative ratings. We tie ratings to actual interactions to force users to also deposit their negative ratings at the reputation server. Otherwise they would not be able to interact any more. Additionally we enhance users' anonymity by limiting timing attacks through the use of transferable-eCash-based payment systems.

5. A. Lorenz, K. Borcea-Pfitzmann, "Role Management in a Privacy-Enhanced Collaborative Environment," in *IADIS International Conference e-Society 2010* [LBP10].
Abstract. Nowadays, social software is in demand in very different settings. So, managing relationships (e.g., social networking sites) and sharing content (e.g., photo sharing), but also collaborative working via the Internet became a widely accepted part of the social lives of people. Especially, collaborative environments provide platforms supporting users in collaboratively creating new ideas, material, and conducting discussions, but also in representing themselves by allowing for according profile management etc. (cf. Koch & Richter 2007). Supporting the users'

privacy in such interactive environments stands in sharp contrast to the objectives of collaboration. However, previous work has shown that different approaches may overcome this ostensible contradiction. One further approach is subject of this paper and consists of a differentiated role management. Accordingly, this paper describes the particular settings of applications shaping Privacy-Enhanced Collaborative Environments (PECE), for which a comprehensive role management has to be realized. The paper discusses the implications on the role concept resulting from the privacy-related settings and introduces a three-dimensional approach for roles in a collaborative environment.

6. B. Kellermann, R. Böhme, "Privacy-Enhanced Event Scheduling," in *Proc. of IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT 2009)* [KB09].
   Abstract. Event schedulers, well-known from groupware and social software, typically share the problem that they disclose detailed availability patterns of their users. This paper distinguishes event scheduling from electronic voting and proposes a privacy-enhanced event scheduling scheme. Based on superposed sending and Diffie–Hellman key agreement, it is designed to be efficient enough for practical implementations while requiring minimal trust in a central entity. Protocols to enable dynamic joining and leaving of participants are given.

7. B. Kellermann, "Datenschutzfreundliche Terminplanung," in *Proc. of the 26th Chaos Communication Congress* [Kel09].
   Abstract. Terminplanungstools sind vielen vielleicht aus Groupwaresystemen bekannt. Ein in letzter Zeit sehr beliebtes Beispiel, welches Terminplanung als stand-alone Web 2.0-Anwendung umsetzt, ist doodle. Doodle sowie andere Lösungen haben das Problem, dass die Vorlieben bzw. Verfügbarkeitsmuster der an der Terminplanung beteiligten Personen veröffentlicht werden. In diesem Beitrag werden ein Protokoll und eine Anwendung vorgestellt, welche das Problem mit homomorpher Verschlüsselung, eine auch im E-Voting verwendete Technik, vermeiden.

8. H. Hedbom, T. Pulls, P. Hjärtquist, A. Laven. "Adding Secure Transparency Logging to the Prime Core," in *Pre-Proc. of the 5th International Summer School: Privacy and Identity Management for Life* [HPHL09].
   Abstract This paper presents a secure privacy preserving log. These types of logs are useful (if not necessary) when constructing transparency services for privacy enhancement. The solution builds on and extends previous work within the area and tries to address the shortcomings of previous solutions regarding privacy issues.

6.3 Privacy of data (WP 2.3)

1. M. Bezzi, "Expressing privacy metrics as one-symbol information," in *Proc. of the 3rd International Workshop on Privacy and Anonymity in the Information Society (PAIS 2010)* [Bez10].
   Abstract. Organizations often need to release microdata without revealing sensitive information. To this scope, data are anonymized and, to assess the quality

of the process, various privacy metrics have been proposed, such as k-anonymity, l-diversity, and t-closeness. These metrics are able to capture different aspects of the disclosure risk, imposing minimal requirements on the association of an individual with the sensitive attributes. If we want to combine them in a optimization problem, we need a common framework able to express all these privacy conditions. Previous studies proposed the notion of mutual information to measure the different kinds of disclosure risks and the utility, but, since mutual information is an average quantity, it is not able to completely express these conditions on single records. We introduce here the notion of one-symbol information (i.e., the contribution to mutual information by a single record) that allows to express the disclosure risk metrics. We also show, with a simple example, how l-diversity and t-closeness can be represented in terms of two different, but equally acceptable, conditions on the information gain.

2. M. Bezzi, G. Montagnon, V. Salzgeber, S. Trabelsi "Sharing data for public security," in *Proc. of the 5th International Summer School on Privacy and Identity Management for Life* [BMST09].
Abstract. Data sharing is a valuable tool for improving security. It allows integrating information from multiple sources to better identify and respond to global security threats. On the other side, sharing of data is limited by privacy and confidentiality. A possible solution is removing or obfuscating part of the data before release (anonymization), and, to this scope, various masking algorithms have been proposed. However, finding the right balance between privacy and the quality of data is often difficult, and it needs a one calibration of the anonymization process. It includes choosing the 'best' set of masking algorithms and an estimation of the risk in releasing the data. Both these processes are rather complex, especially for non-expert users. In this paper, we illustrate the typical issues in the anonymization process, and introduce a tool for assisting the user in the choice of the set of masking transformations. We also propose a caching system to speed up this process over multiple runs on similar datasets. Although, the current version has limited functionalities, and more extensive testing is needed, it is a first step in the direction of developing a user-friendly support tool for anonymization.

3. V. Ciriani, S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, P. Samarati, "Keep a Few: Outsourcing Data while Maintaining Confidentiality," in *Proc. of the 14th European Symposium On Research In Computer Security (ESORICS 2009)* [CDF+09c].
Abstract. We put forward a novel paradigm for preserving privacy in data outsourcing which departs from encryption. The basic idea behind our proposal is to involve the owner in storing a limited portion of the data, and maintaining all data (either at the owner or at external servers) in the clear. We assume a relational context, where the data to be outsourced is contained in a relational table. We then analyze how the relational table can be fragmented, minimizing the load for the data owner. We propose several metrics and present a general framework capturing all of them, with a corresponding algorithm finding a heuristic solution to a family of NP-hard problems.

4. V. Ciriani, S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, P. Samarati, "Enforcing Confidentiality Constraints on Sensitive Databases with Lightweight Trusted Clients," in *Proc. of the 23rd Annual IFIP WG 11.3 Working Conference on Data and Applications Security (DBSec 2009)* [CDF+09a].
Abstract. Existing approaches for protecting sensitive information stored (outsourced) at external "honest-but-curious" servers are typically based on an overlying layer of encryption that is applied on the whole information, or use a combination of fragmentation and encryption. The computational load imposed by encryption makes such approaches not suitable for scenarios with lightweight clients. In this paper, we address this issue and propose a novel model for enforcing privacy requirements on the outsourced information which departs from encryption. The basic idea of our approach is to store a small portion of the data (just enough to break sensitive associations) on the client, which is trusted being under the data owner control, while storing the remaining information in clear form at the external (honest-but-curious) server. We model the problem and provide a solution for it aiming at minimizing the data stored at the client. We also illustrate the execution of queries on the fragmented information.

5. V. Ciriani, S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, P. Samarati, "Fragmentation Design for Efficient Query Execution over Sensitive Distributed Databases," in *Proc. of the 29th International Conference on Distributed Computing Systems (ICDCS 2009)* [CDF+09b].
Abstract. The balance between privacy and utility is a classical problem with an increasing impact on the design of modern information systems. On the one side it is crucial to ensure that sensitive information is properly protected; on the other side, the impact of protection on the workload must be limited as query efficiency and system performance remain a primary requirement. We address this privacy/efficiency balance proposing an approach that, starting from a flexible definition of confidentiality constraints on a relational schema, applies encryption on information in a parsimonious way and mostly relies on fragmentation to protect sensitive associations among attributes. Fragmentation is guided by workload considerations so to minimize the cost of executing queries over fragments. We discuss the minimization problem when fragmenting data and provide a heuristic approach to its solution.

6. S. Trabelsi, V. Salzgeber, M, Bezzi, G. Montagnon. "Data disclosure risk evaluation," in *Proc. of the 4th International Conference on Risks and Security of Internet and Systems (CRiSIS 2009)* [TSBM09].
Abstract. Many companies have to share various information containing private data without being aware about the threats related to such non controlled disclosure. Therefore we propose a solution to support these companies to evaluate the disclosure risk for all their types of data; by recommending the safest configurations using a smart bootstrapping system.

6.4 Access control for the protection of user-generated data (WP 2.4)

1. C.A. Ardagna, M. Cremonini, S. De Capitani di Vimercati, P. Samarati, "An obfuscation-based approach for protecting location privacy," in *IEEE Transaction on Dependable and Secure Computing* [ACDS10].
   Abstract. The pervasive diffusion of mobile communication devices and the technical improvements of location techniques are fostering the development of new applications that use the physical position of users to offer location-based services for business, social, or informational purposes. In such a context, privacy concerns are increasing and call for sophisticated solutions able to guarantee different levels of location privacy to the users. In this paper, we address this problem and present a solution based on different obfuscation operators that, when used individually or in combination, protect the privacy of the location information of users. We also introduce an adversary model and provide an analysis of the proposed obfuscation operators to evaluate their robustness against adversaries aiming to reverse the obfuscation effects to retrieve a location that better approximates the location of the users. Finally, we present some experimental results that validate our solution.

2. C.A. Ardagna, M. Cremonini, S. De Capitani di Vimercati, and P. Samarati, "Access control in location-based services," in *Privacy in Location Based Applications* [ACDS09].
   Abstract. Recent enhancements in location technologies reliability and precision are fostering the development of a new wave of applications that make use of the location information of users. Such applications introduces new aspects of access control which should be addressed. On the one side, precise location information may play an important role and can be used to develop Location-based Access Control (LBAC) systems that integrate traditional access control mechanisms with conditions based on the physical position of users. On the other side, location information of users can be considered sensitive and access control solutions should be developed to protect it against unauthorized accesses and disclosures. In this chapter, we address these two aspects related to the use and protection of location information, discussing existing solutions, open issues, and some research directions.

3. C.A. Ardagna, C. Braghin, M. Cremonini, "Net privacy," in *Computer And Information Security Handbook* [ABC09].
   Abstract. In recent years, large-scale computer networks have become an essential aspect of our daily computing environment: we often rely on a global information infrastructure for ebusiness activities such as home banking, ATM transactions, or shopping online. One of the main scientific and technological challenges in this setting has been to provide security to individuals that operate in possibly untrusted and unknown environments. However, beside threats directly related with computer intrusions, epidemic diffusion of malwares, and plain frauds conducted online, a more subtle although increasing erosion of individuals' privacy has progressed and multiplied. Such an escalating violation of privacy has some direct

harmful consequences-for example, identity thefts have spread in recent years-and negative effects on the general perception of insecurity that many individuals now experience when dealing with online services. Nevertheless, protecting personal privacy from the many parties-business, government, social, or even criminal-which look over the value that personal information have, is an ancient concern of modern society, now increased by the features of the digital infrastructure. In this chapter, we address the privacy issues in the digital society from different points of view, investigating: i) the different aspects that the notion of privacy covers and the debate that the intricate essence of privacy has stimulated; ii) the most common privacy threats and the possible economic aspects that may influence the way privacy is (and especially is not, at the current status) managed in most of the firms; iii) the efforts, in the Computer Science community, to face privacy threats, especially in the context of mobile and database system; iv) the network-based technologies currently available to provide anonymity when communicating over a public network.

4. C. Blundo, S. Cimato, S. De Capitani di Vimercati, A. De Santis, S. Foresti, S. Paraboschi, P. Samarati, "Efficient key management for enforcing access control in outsourced scenarios," in *Proc. of the 24th IFIP TC-11 International Information Security Conference (SEC 2009)* [BCD+09].
   Abstract. Data outsourcing is emerging today as a successful paradigm allowing individuals and organizations to exploit external servers for storing and distributing data. While trusted to properly manage the data, external servers are often not authorized to read them, therefore requiring data to be encrypted. In such a context, the application of an access control policy requires different data to be encrypted with different keys so to allow the external server to directly enforce access control and support selective dissemination and access. The problem therefore emerges of designing solutions for the efficient management of the encryption policy enforcing access control, with the goal of minimizing the number of keys to be maintained by the system and distributed to users. Since such a problem is NP-hard, we propose a heuristic approach to its solution based on a key derivation graph exploiting the relationships among user groups. We experimentally evaluate the performance of our heuristic solution, comparing it with previous approaches.

5. C. Blundo, S. Cimato, S. De Capitani di Vimercati, A. De Santis, S. Foresti, S. Paraboschi, P. Samarati, "Managing key hierarchies for access control enforcement: Heuristic approaches," in *Computers & Security* [BCD+10].
   Abstract. Data outsourcing is emerging today as a successful paradigm allowing individuals and organizations to resort to external servers for storing their data, and sharing them with others. The main problem of this trend is that sensitive data are stored on a site that is not under the data owner's direct control. This scenario poses a major security problem since often the external server is relied upon for ensuring high availability of the data, but it is not authorized to read them. Data need therefore to be encrypted. In such a context, the application of an access control policy requires different data to be encrypted with different keys so to allow the external server to directly enforce access control and support selective dissemination and access. The problem therefore emerges of designing

solutions for the efficient management of an encryption policy enforcing access control, with the goal of minimizing the number of keys to be maintained by the system and distributed to users. In this paper, we prove that the problem of minimizing the number of keys is NP-hard and present alternative approaches for its solution. We first formulate the minimization problem as an instance of an integer linear programming problem and then propose three different families of heuristics, which are based on a key derivation tree exploiting the relationships among user groups. Finally, we experimentally evaluate the performance of our heuristics, comparing them with previous approaches.

6. V. Di Lecce, M. Calabrese, V. Piuri "An ontology-based approach to human telepresence," in *Proc. of the IEEE International Conference on Computational Intelligence for Measurement Systems and Applications (CIMSA 2009)* [DCP09].
Abstract. Detecting human presence automatically is a challenging task since several environmental parameters may affect the quality and the continuity of detection. Although many techniques have been developed so far in the literature to solve this problem, they generally rely on well-defined operational context. Hence, they are sensitive to uncontrolled variables and unpredicted events. In this work an ontology-based approach to human telepresence detection is presented. Contrarily to classic sensor-driven techniques, a top-down methodology is applied. Starting from a formal description of the problem ontology, a set of high-response rate and low-response rate sensors is employed in a computational model. As a consequence of this model, a multi-sensor equipped device has been experimentally setup to conduct measurements on real scenarios. Experiments have been devised to estimate the robustness of the detection. In particular, some preliminary evaluations related to using a minimal set of chemical sensors are reported.

7. V. Di Lecce, A. Amato, V. Piuri, "Data fusion for user presence identification," in *Proc. of the IEEE International Conference on Computational Intelligence for Measurement Systems and Applications (CIMSA 2009)* [DAP09].
Abstract. Aim of this work is to present a new approach to the problem of user presence monitoring in working environments. Particularly, this work is focused on the evaluation of the presence or absence of a user in front of a terminal. This question is of paramount importance in applications requiring the user's presence e.g. video surveillance systems, control centrals, etc. The authors propose a technique of data fusion using signals from various low cost sensors.

8. S. Short, A. Rota, M.A. Rahaman, "XML secure views using semantic access control," in *Proc. of 1st International Workshop on Data Semantics (DataSem 2010)* [SRR10].
Abstract. The OASIS eXtensible Access Control Language (XACML) provides an interoperable tool for writing and enforcing access control policies based on attributes, i.e. characteristics of the entities that take part to the access, such as subjects or actions. Unfortunately, the attribute based approach starts to show its limits when entities exhibit complex relationships, such as semantic relations, which would be easily captured using ontologies instead of attributes. This paper integrates the XACML attribute model with an OWL ontology and describes

a practical privacy filtering application able to filter out information from XML documents, according to a set of XACML semantic privacy policies.

Bibliography

[ABC+08] M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. Malone-Lee, G. Neven, P. Paillier, and H. Shi. Searchable encryption revisited: Consistency properties, relation to anonymous IBE, and extensions. *Journal of Cryptology*, 21(3):350–391, 2008.

[ABC09] C.A. Ardagna, C. Braghin, and M. Cremonini. Net privacy. In *Computer And Information Security Handbook*. Morgan Kaufmann, 2009. to appear.

[ABG+05] G. Aggarwal, M. Bawa, P. Ganesan, H. Garcia-Molina, K. Kenthapadi, R. Motwani, U. Srivastava, D. Thomas, and Y. Xu. Two can keep a secret: A distributed architecture for secure database services. In *Proc. of the Second Biennial Conference on Innovative Data Systems Research (CIDR 2005)*, Asilomar, CA, USA, January 2005.

[ABN10] M. Abdalla, M. Bellare, and G. Neven. Robust encryption. To appear at 7th IACR Theory of Cryptography Conference (TCC), 2010.

[ACBM08] E. Androulaki, S.G. Choi, S.M. Bellovin, and T. Malkin. Reputation systems for anonymous networks. In *Proc. of the 8th International Symposium on Privacy Enhancing Technologies (PET 2008)*, pages 202–218, Berlin, Heidelberg, 2008. Springer-Verlag.

[ACDS09] C.A. Ardagna, M. Cremonini, S. De Capitani di Vimercati, and P. Samarati. Access control in location-based services. In C. Bettini, S. Jajodia, P. Samarati, and S. Wang, editors, *Privacy in location based applications*. Springer, 2009.

[ACDS10] C.A. Ardagna, M. Cremonini, S. De Capitani di Vimercati, and P. Samarati. An obfuscation-based approach for protecting location privacy. *IEEE Transaction on Dependable and Secure Computing*, 2010. pre-print.

[ACJT00] G. Ateniese, J. Camenisch, M. Joye, and G. Tsudik. A practical and provably secure coalition-resistant group signature scheme. In M. Bellare, editor, *CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 255–270, 2000.

[ADDS05] C.A. Ardagna, E. Damiani, S. De Capitani di Vimercati, and P. Samarati. Towards privacy-enhanced authorization policies and languages. *Lecture Notes in Computer Science*, 3654:16, 2005.

[ADFM06] G. Ateniese, A. De Santis, A.L. Ferrara, and B. Masucci. Provably-secure time-bound hierarchical key assignment schemes. In *Proc. of the 13th ACM Conference on Computer and Communications Cecurity (CCS 2006)*, Alexandria, VA, USA, November 2006.

[AFB05] M.J. Atallah, K.B. Frikken, and M. Blanton. Dynamic and efficient key management for access hierarchies. In *Proc. of the 12th ACM Conference on Computer and Communications Cecurity (CCS 2005)*, Alexandria, VA, USA, November 2005.

[AJK09] E. Al-Shaer, S. Jha, and A. D. Keromytis, editors. *Proceedings of the 2009 ACM Conference on Computer and Communications Security, CCS 2009, Chicago, Illinois, USA, November 9-13, 2009*. ACM, 2009.

[AT83] S. Akl and P. Taylor. Cryptographic solution to a problem of access control in a hierarchy. *ACM Transactions on Computer System*, 1(3):239–248, August 1983.

[BBS04] D. Boneh, X. Boyen, and H. Shacham. Short group signatures. In Matthew Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 41–55, Santa Barbara, CA, USA, August 15–19, 2004. Springer, Berlin, Germany.

[BCC+09] M Belenkiy, J. Camenisch, M. Chase, M. Kohlweiss, A. Lysyanskaya, and H. Shacham. Randomizable proofs and delegatable anonymous credentials. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 108–125. Springer, 2009.

[BCD+09] C. Blundo, S. Cimato, S. De Capitani di Vimercati, A. De Santis, S. Foresti, S. Paraboschi, and P. Samarati. Efficient key management for enforcing access control in outsourced scenarios. In *Proc. of the 24th IFIP TC-11 International Information Security Conference (SEC 2009)*, Cyprus, Greece, May 2009.

[BCD+10] C. Blundo, S. Cimato, S. De Capitani di Vimercati, A. De Santis, S. Foresti, S. Paraboschi, and P. Samarati. Managing key hierarchies for access control enforcement: Heuristic approaches. *Computers & Security*, 2010. (to appear).

[BCGS09] P. Bichsel, J. Camenisch, T. Groß, and V. Shoup. Anonymous credentials on a standard java card. In Al-Shaer et al. [AJK09], pages 600–610.

[BCOP04] D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano. Public key encryption with keyword search. In Christian Cachin and J. Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 506–522. Springer, Berlin, Germany, 2004.

[BD03] L. Barkhuus and A. Dey. Location-based services for mobile telephony: a study of user's privacy concerns. In *Proc. of the 9th IFIP TC13 International Conference on Human-Computer Interaction (INTERACT 2003)*, Zurich, Switzerland, September 2003.

[Bez10] M. Bezzi. Expressing privacy metrics as one-symbol information. In *Proceedings of 3rd International Workshop on Privacy and Anonymity in the Information Society (PAIS'10)*, 2010. Accepted for publication.

[BGO07] A. Barak and O. Gluck-Ofri. Degree and reciprocity of self-disclosure in online forums. *CyberPsychology & Behavior*, 10(3):407–417, 2007.

[BMP+09] P. Bichsel, S. Müller, F.-S. Preiss, D. Sommer, and M. Verdicchio. Security and trust through electronic social network-based interactions. In *CSE (4)*, pages 1002–1007. IEEE Computer Society, 2009.

[BMST09] M. Bezzi, G. Montagnon, V. Salzgeber, and S. Trabelsi. Sharing data for public security. In *Pre-Proc. of the 5th International Summer School on Privacy and Identity Management for Life*, Nice, France, September 2009.

[BMW03] M. Bellare, D. Micciancio, and B. Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 614–629, Warsaw, Poland, May 4–8, 2003. Springer, Berlin, Germany.

[Bra93] S. Brands. Electronic cash systems based on the representation problem in groups of prime order. In *Preproceedings of CRYPTO '93*, pages 26.1–26.15, 1993.

[BS04] D. Boneh and H. Shacham. Group signatures with verifier-local revocation. In Vijayalakshmi Atluri, Birgit Pfitzmann, and Patrick McDaniel, editors, *ACM CCS 04: 11th Conference on Computer and Communications Security*, pages 168–177, Washington D.C., USA, October 25–29, 2004. ACM Press.

[BW06] X. Boyen and B. Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *CRYPTO*, pages 290–307, 2006.

[CCKS07] M. Casassa-Mont, S. Crosta, T. Kriegelstein, and D. Sommer. Architecture v2. PRIME Deliverable D14.2.c, March 2007.

[CCS09] J. Camenisch, N. Chandran, and V. Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In Joux [Jou09], pages 351–368.

[CDD+05] A. Ceselli, E. Damiani, S. De Capitani di Vimercati, S. Jajodia, S. Paraboschi, and P. Samarati. Modeling and assessing inference exposure in encrypted databases. *ACM Transactions on Information and System Security*, 8(1):119–152, February 2005.

[CDF+07] V. Ciriani, S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati. Fragmentation and encryption to enforce privacy in data storage. In *Proc. of the 12th European Symposium On Research In Computer Security (ESORICS 2007)*, Dresden, Germany, September 2007.

[CDF+09a] V. Ciriani, S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati. Enforcing confidentiality constraints on sensitive databases with lightweight trusted clients. In *Proc. of the 23rd Annual IFIP WG 11.3 Working Conference on Data and Applications Security (DBSec 2009)*, Montreal, Quebec, Canada, July 2009.

[CDF+09b] V. Ciriani, S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati. Fragmentation design for efficient query execution over sensitive distributed databases. In *Proc. of the 29th International Conference on Distributed Computing Systems (ICDCS 2009)*, Montreal, Quebec, Canada, June 2009.

[CDF+09c] V. Ciriani, S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati. Keep a few: Outsourcing data while maintaining confidentiality. In *Proc. of the 14th European Symposium On Research In Computer Security (ESORICS 2009)*, Saint Malo, France, September 2009.

[CDFS07] V. Ciriani, S. De Capitani di Vimercati, S. Foresti, and P. Samarati. Microdata protection. In T. Yu and S. Jajodia, editors, *Secure Data Management in Decentralized Systems*. Springer-Verlag, 2007.

[CDN09] J. Camenisch, M. Dubovitskaya, and G. Neven. Oblivious transfer with access control. In Al-Shaer et al. [AJK09], pages 131–140.

[CDN10] J. Camenisch, M. Dubovitskaya, and G. Neven. Unlinkable priced oblivious transfer with rechargeable wallets. To appear at Financial Cryptography 2010, 2010.

[Cha82] D. Chaum. Blind signatures for untraceable payments. In *Advances in Cryptology – CRYPTO '82*, pages 199–203, 1982.

[Cha85] D. Chaum. Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM*, 28(10):1030–1044, October 1985.

[Cha88] D. Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1(1):65–75, January 1988.

[Chi01] Chicago Tribune. *Rental firm uses GPS in speeding fine*, 2001.

[CHL05] J. Camenisch, S. Hohenberger, and A. Lysyanskaya. Compact E-Cash. In *EUROCRYPT*, volume 3494 of LNCS, pages 302–321, 2005.

[CKRS09] J. Camenisch, M. Kohlweiss, A. Rial, and C. Sheedy. Blind and anonymous identity-based encryption and authorised private searches on public key encrypted data. In *Public Key Cryptography, 12th International Workshop*

on Practice and Theory in Public Key Cryptosystems, PKC 2009, Lecture Notes in Computer Science, page 19, Irvine,CA,USA, 2009. Springer-Verlag.

[CKS09] J. Camenisch, M. Kohlweiss, and C. Soriente. An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In *Public Key Cryptography, 12th International Workshop on Practice and Theory in Public Key Cryptosystems, PKC 2009*, Lecture Notes in Computer Science, page 20, Irvine,CA,USA, 2009. Springer-Verlag.

[CKY09] J. Camenisch, A. Kiayias, and M. Yung. On the portability of generalized Schnorr proofs. In Joux [Jou09], pages 425–442.

[CL01] J. Camenisch and A. Lysyanskaya. Efficient non-transferable anonymous multi-show credential system with optional anonymity revocation. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 93–118. Springer Verlag, 2001.

[CL04] J. Camenisch and A. Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 56–72, Santa Barbara, CA, USA, August 15–19, 2004. Springer, Berlin, Germany.

[CRJ02] A. Carell, N. Reiband, and I. Jahnke. Computergestütztes kollaboratives lernen: Die bedeutung von partizipation, wissensintegration und der einfluss von rollen. *Journal Hochschuldidaktik*, 13(2), September 2002.

[CS97] J. Camenisch and M. Stadler. Efficient group signature schemes for large groups. In Burt Kaliski, editor, *CRYPTO '97*, volume 1296 of *Lecture Notes in Computer Science*, pages 410–424. Springer Verlag, 1997.

[CS03] J. Camenisch and V. Shoup. Practical verifiable encryption and decryption of discrete logarithms. In D. Boneh, editor, *Advances in Cryptology — CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 126–144, 2003.

[CS09] J. Camenisch and P. Samarati. First report on mechanisms, February 2009. PrimeLife Deliverable 2.1.1.

[Cv91] D. Chaum and E. van Heyst. Group signatures. In Donald W. Davies, editor, *Advances in Cryptology – EUROCRYPT'91*, volume 547 of *Lecture Notes in Computer Science*, pages 257–265, Brighton, UK, April 8–11, 1991. Springer, Berlin, Germany.

[CvH91] D. Chaum and E. van Heyst. Group signatures. In Donald W. Davies, editor, *EUROCRYPT '91*, volume 547 of *Lecture Notes in Computer Science*, pages 257–265. Springer-Verlag, 1991.

[DAP09] V. Di Lecce, A. Amato, and V. Piuri. Data fusion for user presence identification. In *Proc. of the IEEE International Conference on Computational Intelligence for Measurement Systems and Applications (CIMSA 2009)*, Hong Kong, China, May 2009.

[DB03] T. D'Roza and G. Bilchev. An overview of location-based services. *BT Technology Journal*, 21(1):20–27, January 2003.

[DCP09] V. Di Lecce, M. Calabrese, and V. Piuri. An ontology-based approach to human telepresence. In *Proc. of the IEEE International Conference on Computational Intelligence for Measurement Systems and Applications (CIMSA 2009)*, Hong Kong, China, May 2009.

[DDF+06] E. Damiani, S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati. Selective data encryption in outsourced dynamic environments. In *Proc. of the 2nd International Workshop on Views On Designing Complex Architectures (VODCA 2006)*, Bertinoro, Italy, September 2006.

[DFJ+07] S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati. Over-encryption: Management of access control evolution on outsourced data. In *Proc. of the 33rd International Conference on Very Large Data Bases (VLDB 2007)*, Vienna, Austria, September 2007.

[DFM04] A. De Santis, A.L. Ferrara, and B. Masucci. Cryptographic key assignment schemes for any access control policy. *Information Processing Letters*, 92(4):199–205, November 2004.

[DK06] M. Duckham and L. Kulik. Dynamic & mobile GIS: Investigating change in space and time. In *Location privacy and location-aware computing*. Taylor & Francis, 2006.

[DKD+09] C. Diaz, E. Kosta, H. Dekeyser, M. Kohlweiss, and G. Nigusse. Privacy preserving electronic petitions. *Identity in the Information Society*, 1(1):14, 2009.

[Dör08] N. Döring. Reduced social cues / cues filtered out. In N. C. Krämer, S. Schwan, D. Unz, and M. Suckfüll, editors, *Medienpsychologie. Schlüsselbegriffe und Konzepte*, pages 290–297, Stuttgart, 2008. Kohlhammer.

[DP06] C. Delerablée and D. Pointcheval. Dynamic fully anonymous short group signatures. In Phong Q. Nguyen, editor, *Progress in Cryptology - VIETCRYPT 06: 1st International Conference on Cryptology in Vietnam*, volume 4341 of *Lecture Notes in Computer Science*, pages 193–210, Hanoi, Vietnam, September 25–28, 2006. Springer, Berlin, Germany.

[ENI07] ENISA. Position paper. reputation-based systems: a security analysis. available from `http://www.enisa.europa.eu/doc/pdf/deliverables/enisa_pp_reputation_based_system.pdf` (last visit 16/06/09), 2007.

[Eur95] European Parliament. Directive 1995/46/EC of the European parliament and of the council. Official Journal of the European Union, October 1995.

[Gen09] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 169–178, New York, NY, USA, 2009. ACM.

[Ger71] U. Gerhardt. *Rollenanalyse als kritische Soziologie: Ein konzeptueller Rahmen zur empirischen und methodologischen Begründung einer Theorie der Vergesellschaftung.* Luchterhand, Neuwied, Berlin, 1971.

[GPS08] S. D. Galbraith, K. G. Paterson, and N. P. Smart. Pairings for cryptographers. *Discrete Applied Mathematics*, 156:3113–3121, 2008.

[Her02] S.C. Herring. Computer-mediated communication on the Internet. In *Annual Review of Information Science and Technology*, pages 109–168. asis&t, 2002.

[HIML02] H. Hacigümüs, B. Iyer, S. Mehrotra, and C. Li. Executing SQL over encrypted data in the database-service-provider model. In *Proc. of the 2002 ACM SIGMOD International Conference on Management of Data (SIGMOD 2002)*, Madison, WI, USA, June 2002.

[Hol06] J.E. Holt. Logcrypt: forward security and public verification for secure audit logs. In *Proceedings of the 2006 Australasian workshops on Grid computing and e-research - Volume 54*, volume 167 of *ACM International Conference Proceeding Series*, pages 203–211. Australian Computer Society, 2006.

[HPHL09] H. Hedbom, T. Pulls, P. Hjärtquist, and A. Laven. Adding secure transparency logging to the prime core. In *Pre-Proc. of the 5th International Summer School on Privacy and Identity Management for Life*, Nice, France, September 2009.

[Jou09] A. Joux, editor. *Advances in Cryptology - EUROCRYPT 2009, 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cologne, Germany, April 26-30, 2009. Proceedings*, volume 5479 of *Lecture Notes in Computer Science*. Springer, 2009.

[KB09] B. Kellermann and R. Böhme. Privacy-enhanced event scheduling. In *Proc. of IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT 2009)*, Vancouver, Canada, August 2009.

[Kel09] B. Kellermann. Datenschutzfreundliche Terminplanung. In Matthias 'wetterfrosch' Mehldau, editor, *Proceedings of the 26th Chaos Communication Congress*, pages 207–211, Marktstraße 18, 33602 Bielefeld, December 2009. Chaos Computer Club, Art d'Ameublement.

[Kel10] B. Kellermann. Dudle homepage. `http://dudle.inf.tu-dresden.de`, January 2010.

[KGV83] S. Kirkpatrick, C.D. Gelatt Jr., and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

[Laz09] K. Lazarev. Privacy-Preserving Management of Contextual Reputations. Master's thesis, Technische Universität Dresden, May 2009.

[LBP10] A. Lorenz and K. Borcea-Pfitzmann. Role management in a privacy-enhanced collaborative environment. IADIS International Conference e-Society, 2010.

[Lee04] J-W. Lee. *Location-tracing sparks privacy concerns.* Korea Times, 2004. `http://news.naver.com` [Accessed 9 February 2010].

[Lee05] H. Lee. Behavioral strategies for dealing with flaming in an online forum. *The Sociological Quarterly*, 46:385–403, 2005.

[Lev66] M.J. Levy. *Modernization and the structure of societies; A setting for international affairs.* Princeton University Press, Princeton, N. J., USA, 1966.

[Lin36] R. Linton. *The study of man: An introduction.* Appleton Century Crofts, Inc., Appleton Century Crofts, Inc., 1936.

[Loo08] Loopt. `http://www.loopt.com/about/privacy-security`, December 2008.

[LRSW99] A. Lysyanskaya, R. L. Rivest, A. Sahai, and S. Wolf. Pseudonym systems. In *SAC*, volume 1758 of *LNCS*, pages 184–199, 1999.

[Luh84] N. Luhmann. *Soziale Systeme: Grundriss einer allgemeinen Theorie.* Suhrkamp Verlag, 1984.

[Lys02] A. Lysyanskaya. Signature Schemes and Applications to Cryptographic Protocol Design. PhD thesis, Massachusetts Institute of Technology, 2002.

[Mea34] G.H. Mead. *Mind, Self, and Society: From the Standpoint of a Social Behaviorist.* University of Chicago Press, Chicago, USA, 1934.

[MT09] D. Ma and G. Tsudik. A new approach to secure logging. *ACM Transactions on Storage (TOS)*, 5(1), March 2009.

[MTMS85] S. MacKinnon, P. Taylor, H. Meijer, and S.Akl. An optimal algorithm for assigning cryptographic keys to control access in a hierarchy. *IEEE Transactions on Computers*, 34(9):797–802, September 1985.

[Näf10] M. Näf. Doodle homepage. `http://www.doodle.com`, January 2010.

[NSW09] G. Neven, N. Smart, and B. Warinschi. Hash function requirements for Schnorr signatures. *Journal of Mathematical Cryptology*, 3(1):69–87, 2009.

[Pöt09a] S. Pötzsch. Codesheet for content anaylsis of forum postings, May 2009. http://www1.inf.tu-dresden.de/%7Epoetzsch/dud/codesheet.pdf.

[Pöt09b] S. Pötzsch. Untersuchung des Einflusses von wahrgenommener Privatsphäre und Anonymität auf die Kommunikation in einer Online-Community. In S. Fischer, E. Maehle, and R. Reischuk, editors, *Informatik 2009, Im Fokus das Leben, 28 September - 02 October 2009, Lübeck*, volume 154 of *Lecture Notes in Informatics*, pages 2152 – 2165, Bonn, 2009. Gesellschaft für Informatik.

[Pri06] Privacy Rights Clearinghouse/UCAN. *A Chronology of Data Breaches*, 2006. `http://www.privacyrights.org/ar/ChronDataBreaches.htm`.

[Rhe93] H. Rheingold. *The Virtual Community: Homesteading on the Electronic Frontier*. Perseus Books, 1993.

[San88] R.S. Sandhu. Cryptographic implementation of a tree hierarchy for access control. *Information Processing Letters*, 27(2):95–98, February 1988.

[SC09] S. Schiffner and S. Clauß. Using linkability information to attack mix-based anonymity services. In *Proc. of the 9th Privacy Enhancing Technologies Symposium (PETS 2009)*, Seattle, WA, USA, August 2009.

[SCS09] S. Schiffner, S. Clauß, and S. Steinbrecher. Privacy and liveliness for reputation systems. In *Proc. of 6th European PKI Workshop (EuroPKI 2009)*, Pisa, Italy, September 2009.

[SGM09] S. Steinbrecher, S. Groß, and M. Meichau. Jason: A scalable reputation system for the semantic web. In *Proc. of IFIP International Information Security Conference (SEC 2009)*, Cyprus, Greece, May 2009.

[Sha07] H. Shacham. A cramer-shoup encryption scheme from the linear assumption and from progressively weaker linear variants. Cryptology ePrint Archive, Report 2007/074, 2007. `http://eprint.iacr.org/`.

[SK86] L. Sproull and S. Kiesler. Reducing social context cues: Electronic mail in organizational communications. *Management Science*, 32(11):1492–1512, 1986.

[SK98] B. Schneier and J. Kelsey. Cryptographic support for secure logs on untrusted machines. *Proc. of the 7th USENIX Security Symposium*, January 1998.

[SRR10] S. Short, A. Rota, and M.A. Rahaman. Xml secure views using semantic access control. In *Proc. of the 1st International Workshop on Data Semantics (DataSem 2010)*, Lausanne, Switzerland, March 2010.

[SSA06] S. Sackmann, J. Strüker, and R. Accorsi. Personalization in privacy-aware highly dynamic systems. *Communications of the ACM*, 49(9), September 2006.

[Ste09a] S. Steinbrecher. Enhancing multilateral security in and by reputation systems. In *Proceedings of the IFIP/FIDIS Internet Security and Privacy Summer School, Masaryk University Brno, 1-7 September 2008, to be published by Springer*, volume 298 of *IFIP AICT*, pages 135–150. Springer, 2009.

[Ste09b] S. Steinbrecher. First report on supporting mechanisms, February 2009. PrimeLife Heartbeat 2.2.1.

[Ste09c] S. Steinbrecher. First version of tools on user's supporting mechanisms, August 2009. PrimeLife Heartbeat 2.2.4.

[TSBM09] S. Trabelsi, V. Salzgeber, M. Bezzi, and G. Montagnon. Data disclosure risk evaluation. In *Proc. of the 4th International Conference on Risks and Security of Internet and Systems (CRiSIS 2009)*, Toulose, France, October 2009.

[Vos04] M. Voss. Privacy preserving online reputation systems. In *International Information Security Workshops*, pages 245–260. Kluwer, 2004.

[Wes67] A. Westin. *Privacy and Freedom.* Antheneum, 1967.

[WSLP08] K. Wouters, K. Simoens, D. Lathouwers, and B. Preneel. Secure and privacy-friendly logging for egovernment services. In *Proc. of the 3rd International Conference on Availability, Reliability and Security (ARES 2008)*, Barcelona, Catalonia,ES, March 2008.

[Zhu03] H. Zhu. Some issues of role-based collaboration. In *Proc. of Canadian Conference on Electrical and Computer Engineering (CCECE 2003)*, Montreal, Canada, May 2003.

[Zna65] F. Znaniecki. *Social relations and social roles: The unfinished systematic sociology.* Chandler, San Francisco, USA, 1965.