PRIME
Privacy and Identity Management for Europe

Contract N^o 507591

Title:	*Fourth Annual Research Report*
Author:	*WP 16.1*
Editor:	*René Balzer, Sebastian Clauß, Andreas Pfitzmann*
Reviewers:	*Lexi Pimenidis, Jan Camenisch*
Identifier:	*D16.1.m*
Type:	*Deliverable*
Version:	*1.0*
Date:	*May 31, 2008*
Status:	*Final*
Class:	*Public*

Abstract

PRIME envisions that individuals will be able to interact in the information society in a secure and safe way while retaining control of their privacy. While there already exists a fair number of privacy enhancing technologies allowing to build identity management systems that come surprisingly close to realize this vision, there still are considerable gaps between these technologies and whole areas that are still missing solutions adequately addressing privacy. Indeed, the field of *privacy-enhancing* identity management is still relatively young and many challenging problems still require solutions.

This document reports on the research achieved by PRIME's research work packages WP6.1 and WP[8-12].0 during the fourth project year and describes issues and areas which are open for further research. These work packages were investigating graphical user interfaces, authorization models, cryptography mechanisms, communication infrastructures, and user-side and services-side identity management aiming to advance the state of the art in privacy-enhancing identity management in general and w.r.t. the requirements collected by PRIME in particular. While some of the research results have been incorporated into development of prototypes, others will be taken up within follow-up projects the PRIME partners are involved in.

In the appendix, this document lists the abstracts of all the 81 scientific publications produced by the research work packages WP6.1 and WP[8-12].0 throughout the reporting period.

The PRIME project receives research funding from the Community's Sixth Framework Programme and the Swiss Federal Office for Education and Science.

Members of the PRIME consortium:

International Business Machines of Belgium	Belgium
IBM Research GmbH	Switzerland
Unabhängiges Landeszentrum für Datenschutz	Germany
Technische Universität Dresden	Germany
Deutsche Lufthansa AG	Germany
Katholieke Universiteit Leuven	Belgium
T-Mobile Deutschland GmbH	Germany
Hewlett-Packard Ltd.	United Kingdom
Karlstads Universitet	Sweden
Università degli studi di Milano	Italy
Joint Research Centre	Italy
Centre National de la Recherche Scientifique	France
Johann Wolfgang Goethe Universität Frankfurt	Germany
Chaum LLC	United State of America
Rheinisch-Westfälische Technische Hochschule Aachen	Germany
Institut EURECOM	France
Erasmus Universiteit Rotterdam	The Netherlands
Universiteit von Tilburg	The Netherlands
Fondazione Centro San Raffaele del Monte Tabor	Italy
Swisscom AG	Switzerland

Published PRIME documents.
The following PRIME documents are available at `https://www.prime-project.eu`.

Excerpt of project "Description of work"	03-2004
Project presentation	09-2004
Overview of existing assurance methods	09-2004
Tutorial Version 0	09-2004
Requirements Version 0	10-2004
Architecture Version 0	10-2004
Evaluation of early prototypes	12-2004
HCI guidance and proposals	02-2005
Requirements Version 1	05-2005
Framework Version 1	06-2005
Annual Research Report I	04-2005
White Paper Version 1	07-2005
Tutorial Version 1	07-2005
Architecture Version 1	08-2005
Evaluation of Integrated Prototype V1	10-2005
Evaluation of Initial Application Prototypes	03-2006
General Public Tutorial	02-2006
Annual Research Report II	04-2006
Framework Version 2	07-2006
Advanced Tutorial	02-2007
Architecture Version 2	03-2007
Annual Research Report III	04-2007
Evaluation of Integrated Prototype V2	05-2007
White Paper Version 2	06-2007

The PRIME Deliverable Series

Vision and Objectives of PRIME
Information technologies are becoming pervasive and powerful to the point that the privacy of citizens is now at risk. In the Information Society, individuals need to be able to keep their autonomy and to retain control over their personal information, irrespective of their activities. The widening gap between this vision and current practices on electronic information networks undermines individuals' trust and threatens critical domains like mobility, healthcare, and the exercise of democracy. The goal of PRIME is to close this gap. PRIME develops the PRIME Framework to integrate all technical and non-technical aspects of privacy-enhancing identity management and to show how privacy-enhancing technologies can indeed close this gap. PRIME elicits detailed requirements from legal, social, economic, and application points of view and shows how they can be addressed. PRIME will enable the users to effectively control their private sphere thanks to the PRIME Architecture that orchestrates the different privacy-enhancing technologies, including the human-computer interface. To validate its results, PRIME develops prototypes and conducts experiments with end-users in specific application areas. PRIME advances the state of the art far beyond the objectives of the existing initiatives to address foundational technology, through PRIME research on human-computer interface, ontologies, authorization and cryptology, anonymous communications, and privacy-enhancing identity management systems architecture and assurance methods, taking into account legacy and emerging systems. PRIME raises awareness of privacy-enhancing identity management through its white paper and tutorials, as well as press releases, leaflets, slide presentations, and scientific publications. The following PRIME materials are available from http://www.prime-project.eu:

Introduction to PRIME

- Press releases, leaflets, and slide presentations outline the project objectives, approach, and expected results;
- The PRIME White Paper introduces privacy-enhancing identity management issues and PRIME's vision, solutions, and strategies;
- Tutorials introduce major concepts of privacy-enhancing identity management for use by the software development community and the general public.

PRIME technical materials

- PRIME Framework reviews privacy-enhancing identity management issues, PRIME legal, social, and economic requirements, PRIME concepts and models, and PRIME architecture outline;
- PRIME Requirements analyzes in-depth the legal, social, economic, and application requirements. They comprise of generic requirements, as well as specific, scenario-based requirements of selected application areas including eLearning, location-based services, and airport security controls.
- PRIME Architecture describes in-depth the organization and orchestration of the different privacy-enhancing technologies in a coherent PRIME system;

- Annual research reports review the research results gained in PRIME over the past years, and the research agenda for the subsequent years;
- HCI Guidance provides a comprehensive analysis of the Human-Computer Interface requirements and solutions for privacy-enhancing identity management;
- Assurance methods surveys the existing assurance methods that are relevant to privacy-enhancing identity management;
- Evaluation of prototypes assesses the series of early PRIME technology prototypes from the legal, social, and economic standpoints;
- Scientific publications address results produced in all PRIME-related fields within the scope of the project.

PRIME work plan

PRIME global work plan provides an excerpt of the contract with the European Commission.

Foreword

This is PRIME's fourth annual research report.

It describes open research problems in the problem area of privacy-enhancing identity management PRIME has focussed on, reports what progress PRIME has made on these problems and what problems have been solved by PRIME in the reporting period, and then gives an outlook on research issues remaining open for further research. This report is a standalone document, in particular, it does not assume that the annual research reports for PRIME's first, second and third project year are read before this one.

The document has been produced by WP16.1 and edited by René Balzer, Sebastian Clauß and Andreas Pfitzmann. PRIME partners involved in the work packages WP6.1 and WP[8-12].0 performing basic research have contributed to it.

The following are the main contributors of this document:

- Chapter 1 (Introduction) was written by Jan Camenisch, Sebastian Clauß, abhi shelat, and Andreas Pfitzmann.
- Chapter 2 (Summary) was written by Sebastian Clauß.
- Chapter 3 (HCI Research (WP6.1)) was written by Simone Fischer-Hübner.
- Chapter 4 (Authorization Models Research (WP8.0)) was written by Sabrina De Capitani di Vimercati and Pierangela Samarati.
- Chapter 5 (Cryptographic Mechanisms Research (WP9.0)) was written by Markulf Kohlweiss.
- Chapter 6 (Communication Infrastructure Research (WP10.0)) was written by Lexi Pimenidis.
- Chapter 7 (User-side Identity Management Research (WP11.0)) was written by Sebastian Clauß, Stephen Crane, Marit Hansen, Stefanie Pötzsch and Sandra Steinbrecher.
- Chapter 8 (Services-side Identity Management Research (WP12.0)) was written by Pete Bramhall.
- Chapter 9 (Conclusion) was written by Jan Camenisch and Sebastian Clauß.

Table of Contents

List of Acronyms.

ACDF	Access Control Decision Function
API	Application Programming Interface
BNF	Backus Naur Form
EU	European Union
HCI	Human Computer Interaction
HIBE	Hierarchical Identity Based Encryption
IBE	Identity Based Encryption
IDEMIX	Identity Mixer (Credential System by IBM Zurich Research Laboratory)
IDM	Identity Management
IMA	Identity Management Application
IMS	Identity Management System
ISI	Identity Management System Interface
LBS	Location Based Services
MIX	Packet Mixing Anonymizing Network
OASIS	Organization for the Advancement of Structured Information Standards
OWL	Web Ontology Language
PEKS	Public-Key Encryption with Keyword Search
PII	Personal Identifiable Information
PIR	Private Information Retrieval
PKI	Public Key Infrastructure
PM	Policy Management
PRIME	Privacy and Identity Management for Europe
RDF	Resource Description Framework
RDFS	RDF Schema
REL	Rights Expression Language
RSA	Rivest, Shamir, Adleman
SAML	Security Assertions Markup Language
TCG	Trusted Computing Group
TLS	Transport Layer Security
TPM	Trusted Platform Module
VPN	Virtual Private Network
WP	Work Package
WSN	Wireless Sensor Network
XML	Extensible Markup Language

1 Introduction

In everyday non-electronic life, individuals naturally play different roles such as “family member”, “citizen” or “patient”. Moreover they communicate with several other parties. Typically, when individuals perform a certain role or are participating in a certain communication relationship, they do not reveal all personal data about themselves — rather, they disclose parts of their personal data. In this way, each role or communication relationship could be associated with a *partial identity* of this person. Indeed, properly choosing a role and a (partial) identity within the context of a given social situation is a natural (if learnt) human capability.

The problem with electronic life is that customers, citizens, and consumers are challenged with applying the same skill in the on-line world, in which users have to cope with technical complexity, no forgetfulness as with humans, and rapid change. Because of the resulting disenfranchisement between the user and the mechanically generated world she or he navigates, managing identity becomes a crucial ability that requires precision and hence support by an identity management system.

PRIME’s main goal is the design and development of an identity management system that provides maximal privacy to users. The underlying principle of such a system is that *no party should per se learn any information other than it absolutely needs.* Determining which information is indeed needed depends on the particular application and on business and legal requirements. That, however, is the easy problem to solve. The challenge in facilitating information parsimony is that a good identity management system must meet the following non-trivial criteria:

1. First of all, it must ensure that communication is secure, anonymous (i.e., does not reveal potentially private information such as the user’s IP-address or location to anyone), and correct (i.e., the transmitted information is received only by the intended recipient).

2. Whenever the user needs to provide some information about herself, the system should ensure that no additional information is revealed. For instance, if the user is to reveal that she is of major age, she should not be required to provide her birth date or even her name! Moreover, the party who certifies that a user is of age and the party who verifies the statement should not be able to tell whether they communicated with the same user or with different ones. To achieve this, the system needs to support so-called *anonymous credential systems.*

3. Today’s access control systems typically fully identify the user and then decide whether or not the user is entitled to use the service(s) or information requested. However, a privacy-friendly access control system needs to ensure that the user only reveals the information about herself that is necessary to decide that she is entitled. This requires privacy-enabling policy languages as well as privacy-enabling access control systems.

4. Users need assistance to manage their personal information. This is due in part to the volume of the data and the number of different transactions a user participates in. More importantly, users must painstakingly avoid mistakes since mistakes are almost never forgotten in the on-line world! Thus, the user’s interaction with the systems must be intuitive and easy to understand, with assistance for this management either through automation, if possible, or on-line and contextual help.

5. To maximize trust of the individual in the privacy-enhancing identity management system, the system should be as much under her/his control as possible. Whenever possible,

functionality of the identity management system should be realized by personal devices of this particular individual — at least as an option the individual might choose.

Beyond such a system itself, privacy-enhancing identity management also requires privacy-enabling infrastructure services such as identity brokers, traffic anonymizers, and all kinds of certification authorities.

PRIME is building a system that, for the first time, will meet all these requirements: PRIME has a number of work packages that implement state-of-the-art technologies providing the required functionality and that integrate them into a single prototypical system. This prototyping is supported by research work packages which elaborate on these technologies such that

- they can indeed be integrated with each other and unfold their full potential,
- they meet the requirements from a legal, social, economic, and application point of view.

PRIME has six such research work packages, i.e., WP6.1 and WP[8-12].0. The first one, WP6.1 considers human-computer interfaces with special emphasis on privacy and security-relevant interactions. The second, WP8.0, studies authorization models and languages and addresses Criteria 2. The third, WP9.0, addresses Criteria 1 and 2 and researches cryptographic mechanisms including anonymous credential systems. The fourth, WP10.0, is concerned with communication infrastructures to address Criteria 1. The fifth, WP11.0, works on user-side IDM issues such as user interfaces and tools to support the user in managing her identities to address Criteria 4 and 5. Finally, the sixth one, WP12.0, is concerned with enabling and supporting user-centric privacy-friendly identity management from the services side, thereby also addressing Criteria 4 and 5.

To adapt known technologies and to develop new ones in order to meet PRIME's requirements, these research work packages obtain input and feedback from the work packages implementing the (integrated) prototype, from the work package using that integrated prototype to realize application scenarios (i.e., application prototypes), as well as from the various work packages evaluating these technologies and collecting requirements from a legal, social, economic, and application point of view.

The research performed in the fourth project year is therefore based on the research results reached within year three, but input has also been received from legal, social, economic and application requirements, in particular for the e-Learning and Location-Based-Services applications areas. The research work packages have also received the evaluation reports on the development of the PRIME integrated prototype Version 2 and the definition of PRIME Architecture V2 to further guide their research agendas.

2 Summary

This report summarizes the achievements of PRIME's six research work packages in the fourth project year and discusses issues and areas which are open for further research beyond PRIME.

Some research problems (e.g., anonymous communication and location based services) have been investigated by more than one work package. Thereby, each work package has focussed on specific issues of the problem. So, each of these issues is discussed in the section of the work package which has performed research on this particular issue.

Below, we give a brief research summary for each of the chapters. The full report for each work package follows in a separate chapter.

WP6: HCI Research This group has researched user interfaces which are intuitive and intelligible and which also enforce and promote legal privacy principles. Within the reporting period, research focussed on approaches for a simplified privacy policy management for end users, which included research on suitable choices for predefined preference settings solving the privacy interests of most end users. UI approaches for presenting and semi-automatically generating privacy preferences have been developed and evaluated by pilot user tests.

WP8: Authorization Models Recent results from this group include research on various methods of location privacy protection with regards to location based services for mobile communication. Further, research has been done regarding selective data protection in case of outsourced data storage. Finally, research has been done regarding trust management within database management systems.

WP9: Cryptography The cryptography work package has made progress in four areas: (1) improving both efficiency and security of anonymous credential systems, (2) revisiting and scrutinizing statistical disclosure attacks on anonymous communication systems, (3) research on new anonymity related definitions for cryptographic primitives, (4) cryptographic protocols for private service access with regards to location based services.

WP10: Communication Infrastructure This group has researched on different issues regarding anonymous communication networks. Research focused on (1) performance evaluation and improvement of network layer anonymization, (2) finding limits of practical anonymization techniques regarding a wide variety of attacks, (3) research on holistic anonymity, i.e. anonymity when analyzing traffic data (network layer) and credential information (application layer), and (4) Research on law enforcement in spite of usage of anonymization techniques.

WP11: User-Side Identity Management This group has produced research results in five areas: (1) quantification of privacy in terms of anonymity and unlinkability while performing actions supported by privacy-enhancing identity management, (2) research on privacy-respecting design options for reputation systems for future internet communities, (3) user-controlled privacy-enhancing identity management for multilateral interactions, (4) research on issues arising regarding implementation of a third-party information service supporting users in privacy-friendly management of digital identities, and (5) proof-of-concept studies regarding establishing user's trust in enterprises by utilizing a reputation system.

WP12: Server-Side Identity Management This research group has (1) further developed models and prototypes for scalable obligation management, (2) worked towards creating a working prototype for identity-aware devices, (3) specified models, protocols and developed prototypes for various aspects of assurance control, and performed research regarding (4) privacy-compliant access control and (5) private information retrieval to be used for anonymous communication.

3 HCI Research (WP6.1)

3.1 Introduction

PRIME will only be successful if its technologies are accepted and applied by the end users. For this reason, the research and development of user interfaces for PRIME technologies, which are intelligible, user-friendly while compliant with legal privacy principles, and which are mediating trust, play an important role. Such user interfaces have to meet challenges such as:

- **Mediate complex PET concepts:** PRIME and other privacy-enhancing technologies (PETs) are based on technical concepts or constructs such as *pseudonyms*, *unlinkability*, *anonymous credentials* as well as policy negotiation and management that are unfamiliar to many end users and often do not fit to their mental picture of what is technically feasible. Informational self-determination means that users are able to decide how their personal data are used. This should not necessarily have to involve determining how technicalities such as pseudonymization are carried out. From a usability perspective such technicalities should on the contrary rather be invisible to the users. However, when it comes to understanding the risk of being identified across different interactions with one or several service providers, some sort of notion about digital identity must be understood by the user.

- **Provide Security:** PRIME user interfaces also need to be "secure" in the sense that it should have reasonable countermeasures against common types of Internet fraud attacks, such as phishing and spoofing.

- **Being informative while user-friendly:** Users must be well informed about the consequences when releasing data, and consequently there are legal requirements for providing information to the users (e.g. Art.10, 11 EU Directive 95/46/EC) that need to be met by the PRIME user interfaces. Transparency for individuals of personal data processing at the data controller's site is a basic privacy principle and prerequisite for user control. Nevertheless, users should not be confronted with excessive or badly structured information that are usually perceived as bothersome and ignored by the users.

- **Overcome the end user's lack of trust:** Usability tests of early PRIME prototypes have shown that there are problems to make people trust the claims about the privacy enhancing features of the systems (see D6.1.b, Pettersson et al. in [PFHND^{+}05]). Similar findings of a lack of trust were also recently reported by Günther et al. [GS05] in a study on the perception of user control with privacy-enhancing identity management solutions for RFID environments, even though the test users considered the PETs in this study fairly easy to use.

Within the first two project years, the HCI research within PRIME had focussed on researching and developing HCI paradigms, constructs and guidelines for designing intuitive, intelligible and secure user interfaces for privacy-enhancing identity management that enforce legal privacy principles, and that are thereby meeting the first three challenges listed above. In the second and third project year, we have also started to address the last two challenges by conducting research on HCI designs and constructs for user-friendly end user transparency tools as well as research on how to mediate reliable trust and assurance to end users.

The emphasis of the HCI research in PRIME during the last project year has been on researching approaches for a simplified privacy policy management for end users, which partly addresses the first challenge of Mediating complex PET concepts listed above. This has included research on suitable choices of predefined preference settings providing the most privacy-friendly options and protecting the privacy interests of most end users. We also investigated UI approaches for presenting privacy preferences to users and enabling them to semi-automatically generate/customize these privacy preferences "on the fly" when a service's side has been contacted and is requesting personal data.

While progress in the area of simplified policy management has been made, there are still some open issues that need attention as well as further improvements and extensions that could be made. This includes more research on how the user should be best supported in his decision making in case of a mismatch between his/her privacy preferences and a web site's privacy policy. Also, suitable icons for predefined "standard" privacy preferences and privacy policies should be further researched.

3.2 Research Issues and Results

In this section, we describe in more detail our research results and the issues that we have addressed in the area of simplifying policy management for end users. For this, we will present a set of predefined privacy preferences, from which users can choose from, and then discuss UI approaches that allow to customize privacy preferences semi-automatically "on the fly" when a service's side has been contacted and is requesting personal data. Those pre-defined privacy preferences should represent the privacy interests that users might have for various applications using basic settings for managing most of the identity management tasks and should include the most privacy-friendly options.

3.2.1 Background

Several authors have noticed the difficulties for end-users to set security parameters while heavily simplified setting functions do not provide an adequate set of security levels (Krüger, 1999; Whitten & Tygar [WT99], Jendricke & Gerd tom Markotten [JG00], Gerd tom Markotten [Ger02], several works by Steven Furnell, i.a. [Fur04b, Fur04a, Fur05, FJK06], Cranor & Garfinkel [CG05]). Nielsen's report "User Education Is Not the Answer to Security Problems" the accountability of security cannot be the users' responsibility [Nie04]. He adheres to common recommendations about making security a built-in feature of all computing elements and turning on all security settings by default "since most people don't mess with defaults. Then, make it easy to modify settings so that users can get trusted things done without having to open a wide hole for everybody". This has basically also been our approach in PRIME that we have tried to follow.

The P3P privacy bird provides some predefined P3P preference settings, which can be customized by the user during the installation process and via the privacy-bird menu. However, in contrast to the approach that we have taken, P3P does not permit to define more fine-grained privacy preferences that could for instance conditioned on individual data controllers and data values. The privacy bird also does not allow to change privacy preference settings semi-automatically "on the fly". Hence, it is not surprising that a Privacy Bird User Study reported that while those users who changed their privacy settings reported it was relatively easy to do so, only a minority reported changing them several times [CAG02].

3.2.2 Simplified Policy Handling

In PRIME, the user's release policy (or his/her so-called "privacy preferences") defines the user's preferences regarding the release/disclosure of his/her data. At the service's side, a so-called data handling policy (or simply "privacy policy") specifies how and what data are used/processed by the service's side. If personal data are requested from a user by a service's side, the PRIME user-side system can compare ("match") the service's side's privacy policy with the user's release policy (privacy preferences) and warn the user in case of a mismatch. For ordinary users defining and adapting a privacy-friendly release policy is a complex and error-prone task which usually requires some expertise about basic legal privacy concepts and principles. In the non-electronic world no equivalent task exist, which means that ordinary users have no experiences with how to define and manage their release polices. Without assistance, most users would not define and use release policies at all or could accidentally define or choose a release policy, which are not as privacy-friendly as they would like them to have.

In PRIME, we have therefore derived a set of four predefined "standard" privacy preferences, from which a user could choose from and which he/she can fill in with concrete data values or which he/she could customize "on the fly" and store under a new name. The *predefined privacy preferences* (so-called "PrivPrefs") define what types of data may be released for what specific purposes under what specific conditions. In addition to those settings which will be compared with the privacy policies of service's sides when they request personal data from the user, our privacy preferences also to set the type of pseudonymity/level of linkability to be used. This set of predefined privacy preferences should represent the users' privacy interests therefore also includes the most privacy-friendly options for acting anonymously or for releasing as little information as needed for a certain service.

More precisely, the following PrivPrefs have been defined:

The first one is called *"PRIME-Anonymous"*, which should be activated by default if no other PrivPref has been chosen by the user, and is for example useful for anonymous browsing. With the PrivPref *PRIME-Anonymous*, no personally identifiable data is actively released by default. Transaction pseudonyms are used, i.e. user actions should not be linkable beyond the transaction.

Another one is *"Returning Visitor"*, which is for example useful if the user does not want to directly release personally identifiable data, but would like to allow service's sides to store settings, which can then be utilized for later visits. With the PrivPref *Returning Visitor*, no personal data that are directly identifying the user are released. What might however be released are for instance data about personal settings. Besides, visits to the same side are linkable through the use of role-relationship pseudonyms.

The first two PrivPrefs were designed for applications, where personal data are not directly requested from the user. However, for many applications such as e-Shopping, e-Health or e-Government applications, users usually have to provide personal data. For such applications the most privacy-friendly data release policy will be one, which reveals only the minimal amount of data needed for providing the requested services, where the data will only be retained until the services are completed and will not be forwarded to other third parties. Besides, different transaction pseudonymous should be used for different transactions.

The question of what the minimal amount of data is varies between different applications/services and is dependent on the purposes for which the applications will need to collect and process personal data. Hence, we have to define specific PrivPrefs for specific applications. As an example, we defined the PrivPref called *"Minimal Shopping"* for an e-Shopping service, where the customer would like to release only the minimal amount of data needed for this service, which should be retained only until the service is completed. For providing an e-Shopping service, different personal data items will be needed for the purposes of the sub tasks Registering/Placing an order, Delivery (physical or electronic) and Payment. Today, e-Shopping sites are usually not only collecting data about the placed orders from the customers, but are also requesting payment data and address data from their customers, which they then forward to the payment providers and delivery services which are cooperating with them. However, there is usually no need for the e-Shop vendor to know the customer's address or payment information. In [Ber08], we describe a more privacy-enhanced solution, in which the e-Shop requests only the data needed for placing the order, whereas payment details (e.g., credit card details) are requested directly by the payment provider and address details are directly requested by the delivery service. This means that in such a solution personal data is not forwarded to other third parties, but instead directly requested by the parties that need to process these data.

Purpose	**Data types**	**Comments**
Order	Ordered items	For shopping cart
	Session pseudonyms	= session cookie
Physical delivery	Name	Alternative 1
	Address (full)	
	Pin code (received from service prov.)	*Alternative 2*
	Pick up point	
Electronic delivery	Email address	Alternative 1
	Internet link (user is given a link)	*Alternative 2*
Payment	Credit card info	The alternatives here are not mutually exclusive
	Bank account info	
	(anonymous) eCoins	
	Bonus points	
Registration	User name (automatically generated)	This is the minimal need
	Password (automatically generated)	
Marketing	Email address	
	Telephone number	
	Name (for physical contact)	
	Address (in combination with name)	
Commercialization of data	cf. Marketing	"Transfer" contact info
	PRIME excludes profiling data	
Profiling	Ordered items	Here, the user accepts profiling
	User name / user's pseudonym	
	(possibly more data types)	

Table 1: Data types (tentatively suggested) in relation to stated purposes

Table 1 lists the data types that are typically needed for the purposes of the e-Shopping sub tasks. Our *Minimal Shopping* PrivPref only allows data collection for the purposes Order Registration, Delivery (physical or electronic) and Payment, and restricts the type of data to be collected to those listed for those purposes in Table 1 (which can be assumed to be the minimal amount of data needed for those purposes). If more types of data are requested for those or other purposes (as stated in the service's sides' policies), and if the user has chosen the *Minimal Shopping* PrivPref, he/she will be warned that more data are requested than needed. Besides, the *Minimal Shopping* includes the preference setting that personal data should not be forwarded to other third parties, which means that the users will also be warned if this the service's side's privacy policy allows such data transfers.

Finally, we have also pre-defined a PrivPref called *"Profiled Shopping"*, which could be chosen by users who agree to release more data than needed for the primary e-Shopping service, usually in return to other benefits, such as bonus points. With this PrivPref, the user would also agree to release his/her address details for the purpose "Marketing" and would also agree that the Shop could process information about his orders for the purpose "Profiling" as specified in Table 1.

As mentioned above, the PrivPref *PRIME-Anonymous* is activated by default, if no other PrivPref has been chosen by the user. The user should have the possibilities to select another PrivPref before or when contacting a service's side (for this case we will discuss UI approaches in the next section) or after having contacted a side. The predefined PrivPrefs *Minimal Shopping* and *Profiled Shopping* are from the start only defining what data types (rather than concrete data values) may be released for what purposes. For a simplified handling of PrivPrefs, it should however be possible to customize the PrivPrefs and fill in concrete data values "on the fly" rather than demanding that the user fills in the values by hand before he can use those PrivPrefs. This means when a service's side is requesting personal data and the user fills in data values in a form, such as the "Send Personal Data?" dialogue form (see below), he/she will be requested whether he/she would like to save these data values in the PrivPref that is currently activated. In order to guide the user through these phases of PrivPref selection, data collection by usually different parties such as an e-Shop, payment provider and delivery service and PrivPref customization, a wizard-based user interface approach has been developed. It presents a sequence of decision requests according to the privacy preferences of the end user to compile a finally instantiated privacy policy. A short pilot user test confirmed that users percept the splitting into subtasks as simplifying the process and increasing the transparency. Further user tests should examine this statement in more detail. As shown in Figure 1 the assistant informs the user about the overall procedure. It collects all the required personal data, shows the dedicated purpose of the data request and allows to walk through the different stages to check the settings, made before. It possibly shows also available information about the service provider (e.g. seals or reputation data), about the requested certificates as well as the data handling policies and obligations. In our example the dialogue contains sections with the statements about data recipient, stated purposes and required personal data. The wizard

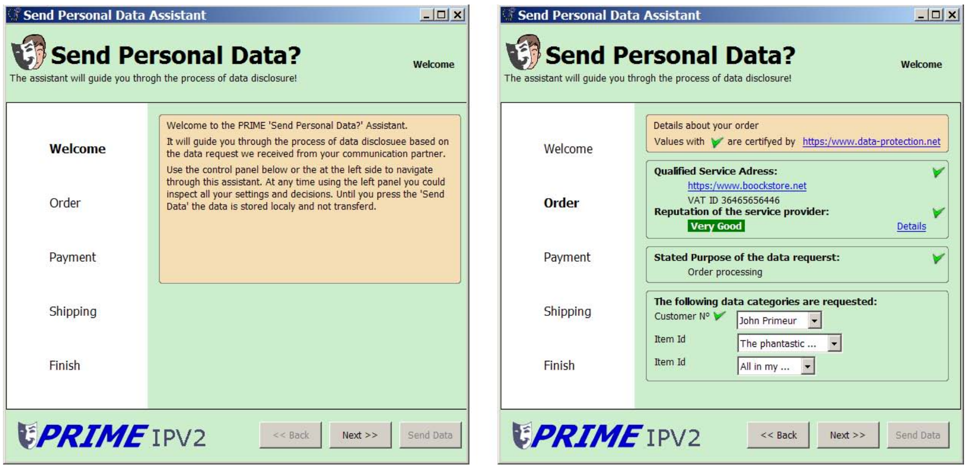

The assistant will guide you throgh the process of data disclosure!

Welcome to the PRIME 'Send Personal Data?' Assistant.

It will guide you through the process of data disclosuee based on the data request we received from your communication partner.

Use the control panel below or the at the left side to navigate through this assistant. At any time using the left panel you could inspect all your settings and decisions. Until you press the 'Send Data' the data is stored localy and not transferd.

The assistant will guide you throgh the process of data disclosure!

Details about your order
Values with ✓ are certifyed by https:/www.data-protection.net

Qualified Service Adress: ✓
https:/www.boockstore.net
VAT ID 36465656446
Reputation of the service provider: ✓
Very Good Details

Stated Purpose of the data requerst: ✓
Order processing

The following data categories are requested:
Customer N° ✓ John Primeur
Item Id The phantastic ...
Item Id All in my ...

Figure 1: An example "Send Personal Data" Assistant

in Figure 1 receives the dedicated data request provided by the service provider.

In comparison to the management of privacy preferences by the P3P privacy bird or other P3P user agents, our PrivPref approach has particularly the following advantages:

- Different PrivPrefs can be defined for different service's sides, whereas P3P only allows the user to define one privacy preference setting which then applies for all web sites that the user visits. Besides, the PrivPrefs allow defining also preference settings on the granularity of concrete data values. Hence, the user can define more fine-grained privacy preferences which provides him/her better protection and may suit better their demands;
- PrivPrefs can be changed and filled in with data values semi-automatically "on the fly", which simplifies the process of changing and customising privacy preferences;
- The predefined PrivPrefs allow to check whether a service's side's privacy policy is conformant with the privacy principle of data minimization and inform users if more data is requested than needed;
- The PrivPrefs allow also to set preferences concerning the type of pseudonymity to be used when using these PrivPrefs.

3.3 Remaining Open Problems

In the area of Simplifying Policy Management for end users, there remain several issues that can be addressed in future:

- Re-investigating alternative UI paradigms for presenting privacy preferences: In the first two project years of PRIME, we have bundled preference settings for personal data and pseudonym types as so-called roles or areas in three main UI paradigms, namely the role-based, the bookmark-based and the TownMap-based paradigms (as reported in the PRIME research report of the second project year). The first two paradigms are traditionally styled while the third one is based on the metaphor of a town map and is an attempt to make preference settings more accessible and, hopefully, understandable to users. On the other hand, the two latter ones share a common approach to the use of preference settings, namely that the selection among the different preference settings (roles and areas, respectively) is implicit when connecting to each service provider. We had researched these UI paradigms before we started our work of PrivPrefs and it could be now re-investigated how suitable the bookmark-based and TownMap-like UI paradigms are for the process of PrivPref selection;
- Further research is needed on how a user should be best informed and supported in his/here decision making in case of a mismatch between his/her privacy preferences (PrivPrefs) and a web site's privacy policy. Our first approach of simply warning the user about the mismatch, resulted in the phenomena that some users customized their privacy preferences, so that they were conforming with the web site's privacy policy in order to avoid further warnings. This however means that they changed their privacy preferences to be less privacy-friendly, which is however the opposite effect than the one which we aimed to achieve;
- Options of using standardized icons for presenting standardized privacy preferences, policies or policy elements should be further researched.

4 Authorization Models Research (WP8.0)

4.1 Introduction

Today's digital business processes increasingly rely on services accessed via a variety of mobile devices and across multiple communication channels [ACD+07d]. Also, terminal devices are now equipped with sensors capable of collecting information from the environment, such as geographical positioning systems, providing a rich context representation regarding both users and the resources they access. This representation includes potentially sensitive personal information, such as user's intended actions, geographical location, and past preferences. While collecting and exploiting rich context data is indeed essential for customizing network-based processes and services, it is well known that context records can be misused well beyond the original intention of their owners. Indeed, personal information is often disclosed to third parties without the consent of legitimate data owners; also, professional services exist specializing on gathering and correlating data from heterogeneous repositories, which permit to build user profiles disclosing sensitive information not voluntarily released by their owners.

In the past few years, increasing awareness of the privacy risks of unauthorized user profiling has led to stricter regulations on personal data storage and sharing. It is now widely acknowledged that business processes requiring large-scale information sharing will become widespread only if their users have some convincing assurance that, while they release the information needed to access a service, disclosure of really sensitive data is not a risk. Unfortunately, some of the emerging technological and organizational requirements for preserving users' privacy are still not completely understood; as a consequence, personal data is often poorly managed and sometimes abused. In general, protecting privacy requires the investigation of different aspects, including the following [DS07].

- *Location information protection* to avoid unauthorized leaks that may cause loss of privacy, for example, on the user's whereabouts. In fact, the pervasive diffusion of mobile communication devices and technical improvements of location technologies are fostering the development of a new wave of applications that use the physical position of individuals to offer location-based services for business, social, or informational purposes [ACD+07d]. In such a context, privacy concerns are increasing and call for sophisticated solutions able to guarantee different levels of location privacy to the users.

- *Data protection* while data are being stored, either on the client side or, more important, on the server side. Therefore, techniques should be adopted both for limiting the possibility of identifying users [DFJ+07c] and for protecting sensitive information about users. These aspects are attracting increasing attention from regulatory bodies and final users, and should be addressed.

- *Security and privacy specifications* to identify under which conditions a party can trust others for their security and privacy. Trust models are one of the techniques be evaluated [DJPS07]. In particular, *digital certificates* (statements certified by given entities) can be used to establish properties of their holder (such as identity, accreditation, or authorizations).

These issues pose several new challenges to the design and implementation of privacy-aware systems. For instance, as far as mobile devices systems are concerned, a major concern is onboard memory and storage limitations. Lightweight terminals require usage logs to be held by the infrastructure, making inference and linking attacks more likely. On the other hand, usage logs need to contain enough information to enable analysis for detection of violations to the privacy policies in place. Another challenge relates to the fact that client and servers alike will not be under the control of authorities trusting each other. Each device and operating system must provide measures to protect the integrity and confidentiality of sensitive personal data and of the privacy control policies. Finally, the lack of resources available on portable devices such as cell phones and laptops may pose some constraints on the effectiveness of purely cryptographic approaches to privacy solutions, adversaries trying to access personal data could have much more computational resources at their disposal than legitimate clients.

In the following, we first describe the advancement status of the research work done in the fourth year of PRIME, which was mainly focused on the privacy issues above-mentioned. We then outline remaining open problems.

4.2 Research Results of the Fourth Project Year

For each research topic addressed, we present the main results.

4.2.1 Location Privacy

The widespread adoption of mobile communication devices combined with technical improvements of location technologies are fostering the development of a new wave of applications that manage physical positions of individuals to offer location-based services for business, social or informational purposes. As an effect of such innovative services, however, privacy concerns are increasing, calling for more sophisticated solutions for providing users with different and manageable levels of privacy [ACDS07c]. Threats to personal privacy in fact are ramping up, as witnessed by recent security incidents targeting privacy of individuals, revealed faulty data management practices, and unauthorized trading of users personal information (including ID thefts and unauthorized profiling). Location information is not immune from such threats and presents new dangers such as stalking or physical harassment [ACD+07d, ACDS07b, ACDS07a].

We address the problem of protecting location privacy of the users by providing a comprehensive solution aimed at preserving location privacy of individuals through artificial perturbations of location information collected by sensing technologies [ACD+07c, ACDS07b, ACD+07a, ACD+07b]. In particular, location information of users is managed by a trusted middleware [ACD+07d, ACDS07a], which enforces user's privacy preferences through obfuscation-based techniques. Our work considers the challenging issue of balancing privacy needs of the users and the need of accuracy requested by location-based service providers to release high-quality services to the users themselves. Our solution introduces the concept of relevance as the dimensionless metric for the location accuracy [ACD+07c, ACDS07b, ACD+07a, ACD+07b]. A relevance value is always associated with locations of users (modelled as circles) and it quantitatively characterizes the degree of privacy artificially introduced into a location measurement. Based on relevance, it is possible to strike a balance between the need of service providers, requiring a certain level of location accuracy, and the need of users, asking to minimize the disclosure of personal location information. Both needs can be expressed as relevance and either quality of online services or location privacy can be adjusted, negotiated or specified as contractual terms. We also provide some obfuscation techniques that are used to degrade the accuracy of users location and to protect their privacy [ACD+07c, ACDS07b, ACD+07a, ACD+07b].

Obfuscation techniques transform an initial area with an initial relevance into an obfuscated area with a final relevance, which satisfies users preferences. In particular, we define the following basic techniques: *i)* the *enlarge* technique degrades the accuracy of an initial location area by enlarging its radius; *ii)* the *shift* technique degrades the accuracy of an initial location area by shifting its center; and *iii)* the *reduce* technique degrades the accuracy of an initial location area by reducing its radius. Given the obfuscation techniques just introduced, users privacy preferences can be satisfied either by using one technique among the three or by composing them [ACD+07a]. Also, we demonstrate that any obfuscation process composed of an arbitrary number of steps can always be reduced to a double obfuscation process, which is composed by a shift technique and an enlarge or reduce technique.

Finally, we observe that the concept of relevance as a metric for location accuracy and privacy is not enough to measure the real privacy protection provided by our obfuscation-based solution. In particular, the degree of robustness of each obfuscation technique must be evaluated with respect to de-obfuscation attempts performed by adversaries.

4.2.2 Location-based Access Control (contributed by JWGFra)

Application of the architecture proposed last year in the different identified scenarios was investigated. The application to mobile devices proved to be feasible. In fact, parts of the location-based service (LBS) access control (AC) implementation were integrated into the PRIME LBS application prototype (V2) via the integrated prototype. In scenarios oriented towards fleet management, we could not identify a requirement for additional privacy functionality, either from the business or legal sides. However, the specifics of this scenario have not been fully investigated.

Progress in LBS AC Integration. The LBS AC components and architecture were integrated (with slight modifications) into the integrated prototype, and have been passed on to at least one application prototype implementation from there. Integration of the architecture in a general-purpose identity management framework has thus been realized, with part of the code even being used towards commercialization.

Progress in LBS AC Design. Based on our experiences with the LBS AC component integration, the LBS AP, and the work on cryptography for LBS done within PRIME, a vision of an LBS AC PET was developed, offering functionality comparable with advanced PETs, while minimizing impact on the infrastructure.

Progress on Economic and Legal Framework for LBS AC. In addition to the technical designs, a technological, economic and legal framework for LBS access control was developed. This already led to several publications [KZSD07, KZSD07, ZFR+07b, ZFR+07a] — more effort could be applied here due to synergies with Activity 4, where there were additional requests raised by project lead. Additionally, this corresponds to the open questions described by last year's report.

4.2.3 Data Protection

The continuous growth of the amount of digital information to be stored and widely distributed, together with the always increasing storage, supports the view that service providers will be more and more requested to be responsible for the storage and the efficient and reliable distribution of content produced by others, realizing a "data outsourcing" architecture on a wide scale. This important trend is particularly clear when we look at the success of services like YouTube, Flickr, Blogger, MySpace, and many others in the "social networking" environment.

When storage and distribution do not involve publicly releasable resources, selective access techniques must be enforced. In this context, it is legitimate for the data owner to demand the data not to be disclosed to the service provider itself, which, while trustworthy to properly carry out the resource distribution functions, should not be allowed access to the resource content. The problem of outsourcing resource management to a "honest but curious" service has recently received considerable attention by the research community and several advancements have been proposed.

In this context, we look at the problem of defining and assigning cryptographic keys to users, by exploiting hierarchical key assignment schemes, and of efficiently supporting policy changes [DFJ+07a, DFJ+07b]. Indeed, in scenarios involving potentially huge sets of resources of considerable size, re-encryption and re-transmission by the owner may not be acceptable. The advantage compared with a solution requiring to re-send a novel encrypted version of the resource is typically huge and arbitrarily large (if the resource has a size of 1 GByte and the request to the server requires a 100-byte packet, in terms of network traffic, compared to the transmission of the re-encrypted resource, the improvement is in the order of 10^7). In particular, the main contributions of our work can be summarized as follows.

- We propose a formal base model for the correct application of selective encryption [DDF+07a]. The model allows the definition of an encryption policy equivalent to the authorization policy to be enforced on the resources. We note that, while it is in principle advisable to leave authorization-based access control and cryptographic protection separate, in the outsourcing scenario such a combination can prove successful: selective encryption allows selective access to be enforced by the service provider itself without the owner intervention.

- Building on the base model, we propose the use of a two-layer approach to enforce selective encryption without requesting the owner to re-encrypt the resources every time there is a change in the authorization policy. The first layer of encryption is applied by the data owner at initialization time (when releasing the resource for outsourcing), the second layer of encryption is applied by the service itself to take care of dynamic policy changes. Intuitively, the two-layer encryption allows the owner to outsource, besides the resource storage and dissemination, the authorization policy management, while not releasing data to the provider.

- We provide a characterization of the different views of the resources by different users and characterize potential risks of information exposures due to dynamic policy changes. The investigation allows us to conclude that, while an exposure risk may exist, it is well defined and identifiable. This allows the owner to address the problem and minimize it at design time. Also, the fact that exposure arises only in specific situations and over well identified resources, allows the owner to completely eliminate it by resorting to re-encryption when necessary.

An important strength of our solution is that it does not substitute the current proposals, rather it complements them, enabling them to support encryption in a selective form and easily enforce dynamic policy changes.

In [CDF+07] we also propose an approach combining fragmentation and encryption that allows storing data on a single server and minimizes the amount of data represented only in encrypted format, therefore allowing for efficient query execution. Indeed, to guarantee the privacy of a collection of data, the current approaches encrypt all the data. The assumption underlying such approaches is that all the data are equally sensitive and therefore encryption is a price to be paid to protect them. This assumption is typically an overkill in many scenarios. As a matter of fact, in many situations data are not sensitive per se; what is sensitive is their association with other data. As a simple example, in a hospital the list of illnesses cured or the list of patients could be made publicly available, while the association of specific illnesses to individual patients is sensitive and must be protected. Hence, there is no need to encrypt both illnesses and patients if there are alternative ways to protect the association between them. We therefore propose an alternative approach that is based on the definition of *confidentiality constraints* as a simple, yet powerful, way to capture privacy requirements. We also provide a model formalizing the application of data fragmentation and encryption, which captures properties related to the correct representation of the data while minimizing encryption and fragmentation. Finally, we propose a heuristic algorithm for the concrete identification of a fragmentation solution that satisfies the properties specified.

4.3 Security and Privacy Specifications

Trust management systems allow possibly unknown parties to establish trust based on (certified) information that each party can present to the counterpart at the time of interaction. The certified information (or certified attributes) is substantiated by digital certificates issued by certification authorities. The authentication phase is managed in a robust way exploiting the benefits of asymmetric cryptography. Servers supporting trust management can then regulate access based on attributes (identities or more general properties) that clients requesting access present. This is often the base on which flexible authorizations can be defined, using certified attributes as parameters in the specification of the resource or subject of an authorization. While several approaches have been proposed for trust management and significant steps have been made in this direction, a major obstacle that still exists in the realization of the benefits of this paradigm is represented by the lack of adequate support in the DBMS.

In [DJPS07], we propose a trust management model for DBMSs. The model identifies and adapts trust management concepts for their handling within relational databases. It is accompanied by a SQL syntax, which allows a seamless integration with existing database systems and demonstrates the high-level abstraction that a database designer can use to apply these concepts. The model introduces a novel mechanism for the retrieval of the delegation certificates, which we demonstrate can be realized in an efficient way. We also illustrate the basic techniques on which a mechanism efficiently enforcing the model within a modern relational engine can be built.

4.4 Remaining Open Problems

Open issues that will be investigated include the following.

Location privacy. The work on location privacy leaves space for further advancements.

- We plan to extend our obfuscation-based solution by considering the possibility of using Gaussian-like distributions and complex location measurement shapes; introducing map constraints in the computation of obfuscated areas; defining additional techniques that protect the privacy of the users by degrading the temporal accuracy of their location measurements.

- We plan to evaluate the robustness of our obfuscation-based solution with respect to possible de-obfuscation attacks made by adversaries. In fact, the definition of an adversary model will allow us to measure the real privacy protection provided by the obfuscation techniques we proposed. This evaluation will take into consideration the fact that an adversary may adopt different strategies and may use contextual information in an attempt to reverse the protection granted by obfuscation.

Location based access control. Due to time constraints, we could not publish the envisioned design of an LBS AC PET offering functionality comparable with advanced PETs, while minimizing impact on the infrastructure. The current LBS AP does not use the fully fledged LBS AC solution, but rather uses components of the LBS AC for acquiring and controlling access to user location information.

Data protection. There are two main directions that will be further investigated.

- We plan to improve our fragmentation and encryption solution by computing a fragmentation that looks carefully at the performance issues and takes into account the profile of the query load on the server.
- We plan to extend the proposed model and data outsourcing access control architecture for supporting flexible applications by integrating them with the Web paradigm, and by defining techniques for the efficient execution of queries.

Security and privacy specifications. Although we believe that our solution that integrates trust management with DBMS authorizations can be immediately implemented by DBMS producers and used by DBAs, the soundness of the proposed model needs to be verified.

- We plan to implement a proof-of-concept prototype to the aim of empirically verify the soundness of our trust management model for relational databases. The prototype will extend the well known PostgreSQL DBMS with our trust management model.

5 Cryptographic Mechanisms Research (WP9.0)

5.1 Introduction

Cryptography provides the core technical mechanisms to achieve, maintain, and protect users' privacy. These mechanisms range from encryption and signature schemes that enable secure and private communication over higher-level protocols such as electronic voting to complex multiparty protocols for secure function evaluation. This activity is concerned with research on such mechanisms, in particular those for privacy-enhancing identity management. The principle goal is to provide and improve on the mechanisms that are needed by PRIME's integrated prototype such a credential systems, onion encryption schemes, or particular mechanisms for the application prototypes (location based services, airport security, and e-Learning). On the other hand, this activity also aims to perform research on privacy enhancing cryptographic mechanisms in general such as group signatures or blind signature schemes.

In the reporting period, WP9.0 has improved the state of the art in four major areas:

Anonymous Credentials We have continued our research on new credential systems to improve both the efficiency [CHP07] and the security of existing solutions [BCKL08]. We also investigated the possibility to delegate credentials anonymously.

Anonymous Communication We revisit and scrutinize statistical disclosure attacks in [DDT07] and [DTD07]. The first looks at how timing information about messages and their replies can help the attacker. Possible counter measures were proposed in [SCM05], and analyzed and improved in [SP07a, PS07] and [SP07b] respectively. The latter points at new directions for anonymity attacks based on *attacker uncertainty* that merges profile information about user with information about the input output behaviour of concrete anonymity systems such as pool mixes.

Definitions of Anonymity We continued to research anonymity related cryptographic definitions. We provide a provable-security treatment of the notion of a 'robust' encryption scheme, namely one where the decryption algorithm rejects when the 'wrong' secret key is used. We investigate applications to auctions, searchable encryption and anonymous wireless communication. We also finished our work on definitions for blind signatures, oblivious transfer [CNas07], and location-based services [KFF$^+$07].

Private Service Access The easiest way to implement electronic services is to exchange all necessary transaction related information, such as policies and personal data, beforehand, and then base business decisions and service provisioning on this information. Doing so may reveal more information than desired by the parties involved.

In the last year we studied the particular case of *location-based service*. In a conventional setup, the location provider simply gives the user's location to the service provider who then forwards the relevant service data to the user. This approach has obvious privacy problems. We proposed improved solutions based on secure and efficient oblivious transfer protocols [CNas07] in which the service provider does not learn the user's location and a new type of *proxy* oblivious transfer protocol in which the communication and computation burden of the protocol is offloaded from the recipient [KFF$^+$07]. We also investigated ways to incorporate privacy friendly service selection and payment protocols into such a system. This approach can be generalized to selling digital goods, and is an example how in the electronic world we can protect privacy better than it appears judging from our experiences in the physical world. Another scenario where we can achieve stronger privacy in the electronic world involves searches on encrypted data.

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

Remaining Open Problems. With the end of the project, research within PRIME will also come to an end. Many challenges, however, remain. Some of them just now becoming apparent, on account of the new findings we made within the course of the last four years. Some of these challenges will be taken up by one of the follow up projects of PRIME, PrimeLife. This follow-up project will be 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, can prove useful within the thick social fabric of today's Internet where private information need not only be hidden but needs to be processed and shared in a responsible manner.

5.2 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 licence to show that she is of age, a credential system allow 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 to make the existing protocols more efficient, realize new features as to enable new application, and to provide schemes that are based on simpler cryptographic assumptions.

Progress in the Fourth Project Year

In the fourth project year, we have achieved two main break through results in the area of anonymous credential systems: (1) efficiency through batch verification, (2) efficient non-interactive schemes provably secure in the standard model.

Batch verification. In order to achieve both small communication cost for credential shows and short verification times, it seems that one needs to improve on the verification of bilinear-map based schemes. In previous work [CHP07] we looked at batch verification of short signatures from potentially different signers on potentially different messages. It showed how to batch verify the Waters identity-based signatures [Wat05] and proposed a new signature scheme (related to CL signatures but not yet usable for anonymous credentials) that batches well.

Non-interactive credentials without ideal assumptions. A CL-signature scheme consists of a signature scheme, a commitment scheme, and (1) an interactive protocol for obtaining a signature on a committed value, and (2) a proof system for proving that the contents of a commitment has been signed.

The proof system used in (2) above is critical for the performance and security of CL-signatures (and thus of a credential system). It is executed every time a credential is shown. CL-signatures are based on so called Σ-proof protocols which are either interactive or rely on the ideal random oracle assumption. Non-interactive proofs are desirable as they reduce the round complexity of a proof system.

In [BCKL08] we give a definition of security for P-signatures that have non-interactive proofs and show how they can be realized under appropriate assumptions about groups with a bilinear map. We make extensive use of the powerful suite of non-interactive proof techniques due to Groth and Sahai [GS08]. Our P-signatures enable, for the first time, the design of a practical non-interactive anonymous credential system whose security does not rely on the random oracle model. In addition, they may serve as a useful building block for other privacy-preserving authentication mechanisms.

Delegable anonymous credentials. 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 obtained first results using the Groth and Sahai proofs [GS08].

5.3 Anonymous Communication

Anonymous communications allow for *anonymity* of the sender, recipient, or both for some information to be transmitted. We survey the existing systems in the field of anonymous communication in [DD08], focusing on such issues as taxonomy of systems, threat models and usage models of theses systems, and the various security and communication property trade-offs made in such systems.

In [DTD07], we challenge the assumption that anonymity (measured in the form of *entropy*) always necessarily decreases given more information, and draw attention to the key factor of as the main metric by which anonymity is governed.

We have also performed extensions on existing attacks, particularly the *statistical disclosure attack* [Dan03]. In [DDT07], we extend the work in [MD04, SCM05] and show again that receives of sent messages are at greater risk to a powerful statistical disclosure attack.

We evaluated several schemes for the revocation of user anonymity, and found them vulnerable to additional privacy risks for existing (legitimate) users as well as subversion of the revocation mechanism itself by misbehaving users, and formalized these new attacks in [DS08].

We also presented a system for evaluating the reliability of mix-nets, where multiple nodes exist and may be in various states of operational capacity at any given time [KPS07]. In the same paper, we discuss the impact of the intersection attack [BPS00] on such reliability services, and propose a system that utilizes a Byzantine agreement protocol to thwart such attacks.

Finally, we evaluated the threat model that existing users of the Tor anonymity system [DMS04] operate under, and suggest that their threat model may not match the protections offered by the Tor system [Sas07] due to the operation of exit nodes by untrusted authorities who are able to monitor all clear-text communication transmitted through these nodes, and potentially steal such information as credit card data, username and password combinations, as well as monitoring sensitive communications. End-to-end encryption solves these problems, but real world analysis of the threat shows that end-to-end encryption is rarely used by the casual user, the same user profile targeted by Tor.

5.4 Anonymity Related Security Definitions for Cryptographic Primitives

Finding the right security definitions for cryptographic primitives is a fundamental and on-going task in the field. We continued our work on definitions for provable secure primitives related to anonymity. The aim of our effort is both to move to a more solid basis for designing such primitives, and to gain insights into how the properties defined in the security definitions of the primitives influence the security of protocols constructed from these primitives.

We provide a provable-security treatment of the notion of a "robust" encryption scheme, namely one where the decryption algorithm rejects when the "wrong" secret key is used. It makes explicit the so-far implicit notion of robustness, defines it formally, and investigates provably achieving it. We find that contrary to intuition, robustness, at least in combination with privacy and anonymity as required by applications, is actually rarely if ever present, and obvious ways to confer it fail. We however provide ways to efficiently confer it without sacrificing other security properties. As a consequence, we obtain several new applications and results, in areas such as auctions, public-key encryption with keyword search (PEKS), and anonymous communication. But we believe that "naming" and provably achieving this so-far under-the-covers notion of robustness is important beyond this, from the point of view of clarifying and helping to fill gaps in the literature, and of making encryption more resistant to misuse.

In the case of blind signatures, we defined *selective-failure blindness*. An attack on this property would allow the signer to force the protocol to fail only for messages "My vote goes to Barack Obama". Affected users could complain that they did not receive a valid signature, but by doing so they give up their privacy. We also noticed a similar problem in existing adaptive oblivious transfer schemes [CNas07]. To counter these problems we defined simulation-based security notions for adaptive oblivious transfer which can be efficiently implemented, and which form the paradigm for our work on location-based services. The task of defining the security requirements for location-based services is also an active area of work [KFF+07].

5.5 Private Service Access

We are interested in protocols that allow users to access services without revealing their access patterns. Such techniques are known as oblivious transfer and Private Information Retrieval (PIR). We work on their definition, their security [CNas07], as well as on their use in privacy enhancing applications [KFF+07, PS07, SP07a, SP07b] (e.g., privately retrieving one's emails from an email server).

The Pynchon Gate system [SCM05], a pseudonym-server based on the mix-net nym-server architecture, and modified to work with information-theoretic PIR, addresses many of the problems with existing *recipient anonymity*. We have identified attacks in that system that can lead to denial of service, allowing a Byzantine attacker to break the utility, but not the anonymity, of the proposed system [SP07a], and presented an initial proposal for a solution [SP07b], using a "cut-and-choose" protocol. We also discussed benign covert channels in the same system [PS07], leaving open the potential for future work using such channels as a primitive for anonymous communication and control systems.

We finished work started in the third project year on location-based services [CNas07, KFF+07]. In today's systems the location information is simply transferred to the location based service which in most cases is unacceptable from a privacy point of view. We investigated cryptographic protocols that offer the same functionality without sacrificing the users privacy. To achieve this goal, we developed new efficient versions of oblivious transfer. We consider an adaptive version of oblivious transfer in which the sender (the location-based service) and receiver first run an initialization phase during which the sender commits to a "database" containing his messages. Later on, the sender and receiver interact as before so that the receiver can retrieve some message for a given location. Now however, we allow the receiver to interact with the sender k additional times, one interaction after the other, in order to retrieve k additional values from the sender's database. Note that in addition to location-based service, this type of adaptive oblivious transfer problem is central to a variety of other practical problems such as "patent searches" and "treasure hunting".

Our protocols meet a strengthened simulation based security definition for oblivious transfer (OT). This is arguably the strongest model of security for OT. The OT primitives are described as an ideal functionality, and the protocol is such that any malicious sender (receiver) has a corresponding ideal-world counterpart (simulator). Here the protocol guarantees security even when run in arbitrarily complicated settings, e.g. concurrently with other protocols.

Moreover, our protocols can be easily transformed into proxy oblivious transfer. In a normal OT, there are two parties, a sender with a list and a receiver with an index in the list to retrieve. In a proxy OT, the index information is split between two parties—a receiver and a receiver-helper. The protocol guarantees receiver-privacy, meaning that not even the receiver-helper learns the full index and the information retrieved, even though the brunt of the protocol communication is handled by the receiver-helper. The proxy OT is quite applicable to the location-based services in which the receiver might be a cell-phone, while the receiver-helper (in this case, the location provider) can have much better communication channels and computational resources.

Furthermore we extended the oblivious transfer protocol with a private service selection protocol, which allows users to subscribe and pay for multiple LBS services at the same time, while hiding from all participants, which services they are actually accessing. In addition a payment solution derived from techniques also used in voting protocols allows the LBS services to be paid according to their popularity, without revealing which service was accessed by which user.

We investigate how anonymous credentials can be combined with oblivious service access protocols to allow for fine grained access control policies on OT messages. Each message of the OT is associated with attributes that the credential of the accessing user needs to match with. Using a restricted oblivious transfer protocol it becomes possible to enforce these conditions without revealing anything about the credential of the user or the message accessed. For instance a message may be associated with a minimum age for users or with a minimum subscription level. For location-based services this allows to associate different geographic areas to subscription types.

We also studied the state-of-the-art of homomorphic cryptosystems in [WDH07]. Homomorphic properties enable operations in the encrypted domain and are useful for various privacy friendly applications, such as secure watermark protocols and anonymous fingerprinting [DP08], private search [DD07] (search in encrypted databases), privacy-preserving data mining, electronic voting, etc.

Buyer-seller watermarking protocols integrate digital watermarking and fingerprinting with cryptography and have applications in the realm of privacy and copyright protection in e-Commerce. In [DP08], we propose two new efficient buyer-seller watermarking protocols with enhanced anonymity control based on group signatures.

In [DD07] we investigate private keyword search, a technique that allows for searching and retrieving documents matching certain keywords without revealing the search criteria. We improve the space efficiency of the Ostrovsky et al. Private Search scheme, by describing methods that require considerably shorter buffers for returning the results of the search.

5.6 Other Relevant Areas

Database privacy. We performed research in the area of database privacy. In [DDF$^+$07b], we investigated negative databases. A negative database is a privacy-preserving storage system that allows to efficiently test if an entry is present, but makes it hard to enumerate all encoded entries. We improve significantly over previous work presented by showing constructions for negative databases reducible to the security of cryptographic hash functions or the hardness of the Discrete-Logarithm problem.

Anonymity factors in deployed hardware. We examined the hardware security model in the OLPC XO laptop system, and uncovered serious privacy and anonymity threats presented as unintentional side-effects of the anti-theft system implemented in these devices [PSC08]. This is preliminary work which will be further explored to present privacy preserving alternatives to such mechanisms. Without such a solution, efforts such as PRIME will not have any impact on users of these devices, or ones inspired by them.

File system privacy. Steganographic file systems allow for the concealment of data without an attacker being aware that data exists in the file system. In [TDDP07], we showed that the continuously-observable steganographic file system proposals, which allow a user to remotely store user files on a raw storage device, fail to meet their security model of allowing plausible deniability in the face of an adversary able to observe traffic to the device and perform traffic analysis against this communication.

6 Communication Infrastructure Research (WP10.0)

In this section we will display the research goals, achievements and remaining open problems of the work package 10.0: Research on the Communication Infrastructure. We show the progress that has been achieved in accordance to the requirements of the complete project. Finally we show future directions of our work.

6.1 Introduction

Typical network traffic in global networks like the Internet is neither confidential nor secure in nearly all aspects: the content of most network connections is not encrypted, the participants are not authenticated, and the identity of the entities is easily discoverable. While there might be exceptions to this, a surprisingly large number of e-Commerce sites doesn't prevent impersonation attacks, sometimes called phishing or man-in-the-middle attacks, or doesn't hide the customer's data from prying eyes while it is transmitted. Also authentication is often done in an insecure way, e.g., by means of the weak passwords chosen by the user. Additionally, many service providers are unable to protect their services and stored data from attackers; in case of an information leak they are also unable to trace and prosecute the offender.

On the other hand, customers are not aware of the risks in the world wide environments and assume that the problems of security and privacy have been solved by the application designers and service providers. This assumption bases on the fact that customers often don't have a technical background and is thus wrong, i.e., the list of issues from above can be continued to any specific extend. It is also a common misunderstanding that a system can be secure or usable, but not both at the same time.

There are actually good solutions for encrypted and authenticated traffic. The most prominent example is TLS, whose basic building blocks are well researched. This includes a secure key exchange (as of Diffie-Hellman), public key infrastructures (e.g., in the format X.509), and even different types of encryption. This makes one wonder, why there are still newly developed programs that don't make use of these tools, i.e., the majority of instant messaging tools transmit the content, which is often of very private nature, in plain. While end users are missing the awareness on their safety, developers find it tedious to make use of the complicated interfaces of the programming libraries needed for encryption and proper authentication.

The same applies even more to the network layer anonymization. These techniques are needed because it is sometimes not enough to hide the content of a confidential network connection, but also there is a need to hide the sender, the recipient, or both. Special care has to be taken to hide this information, as the pure removal or encryption of it will result in the message being dropped on its way. Note that even an innocent looking IP-address might be enough to uniquely determine a person's identity. Since even basic research on this topic is still incomplete, there are only two tools that are mature enough to be candidates for use in productive services: JAP/AN.ON [BFK00] and Tor [DMS04]. In addition to the same issues that encryption suffers from, these products have performance and scalability issues, offer no protection against strong adversaries, and have a connection drop-out rate that is too high for the e-Commerce. As such they are more a type of beta versions, but we can learn valuable lessons from their deployment to develop schemes for anonymized communication.

The goal of WP 10.0 is to devise solutions to those problems:

- *Performance of network layer anonymization* is a hard problem. Current anonymizing networks are client-server based, and not peer-to-peer networks. This leads to scalability problems, as typically the available bandwidth is fixed and lower than the one needed by the continuously growing size of clients. This can be solved by switching to a peer-to-peer network, where all clients contribute some of their bandwidth to the network disposal. On the other hand, switching the paradigm from client-server to peer-to-peer rises a number of new problems. The main three of them, that have to be solved, are:

1. How to find a needed autonomous peer-to-peer network, if there are no dedicated directory servers made publicly available?
2. How to address nodes and route data in the network, whose topology may change rapidly (e.g. every minute) in some segments, while some others remain stable for a longer period of time?
3. How to achieve a good QoS under the given circumstances?

There is a clear need for the new algorithms dealing with dynamic environments, i.e. above described networks.

- *The limits of anonymization techniques* are still unknown. There are only few works that try to estimate the security provided by certain techniques. Without the possibility to measure the quality of a product's security, one is unable to encourage people to use these systems. The degree of anonymity has to be *calculated* and *communicated* to the user so that he is aware of the current situation and can act accordingly. Alternatively, some mechanism in the client's access tool can act on behalf of the user. The level of security that is provided to the user has to be determined either locally at the user side, or with the help of the publicly available information of the network's status. This plays a crucial role in adapting the user behaviour and achieving user acceptance. A correct estimation or calculation of this value is the key to strong anonymity in open environments.

  Anonymization on the traffic layer hides the communication partners' IP address, i.e. supports sender and receiver anonymity. One major problem with mix networks is that the parameters can only be chosen to either support a high level of security, or a high performance, but not both at the same time. Additionally, the security of mix networks depends on the size of the user base. But if the network's latency is getting higher, users will abandon the network, thus lowering the network's protection for the remaining ones. Thus, it's crucial for anonymizing networks to choose the parameters in such a way, that the users are secure enough, and the delay of messages is still acceptable.

  The proposed research direction leads to a deeper understanding on the level of protection granted by anonymizing techniques, given a set of parameters. Today a number of system parameters in mix networks are chosen rather arbitrarily. Besides providing users with a measure of their degree of anonymity and thus raising the user awareness, our methodology could be used to calculate the parameters optimal from the user's points of view, taking into account a trade-off between security and performance.

- *Supporting law enforcement* is crucial to wide deployment and acceptance of network layer anonymization techniques. While massive deployment is a key factor, it will necessarily lead to some abuse. We therefore need to find ways to revoke certain individuals anonymity without giving away the identity of the other users. Additionally it has to be made sure that there is no possibility to abuse this mechanisms.

  Without technically feasible solutions that are widely accepted by the user community it is not improbable that legislative measures would restrict usage of cryptography and anonymization. On the other hand, this would lead to a total leak of privacy for legal users while criminals with specific knowledge and skills could still remain unidentifiable.

A solution to the above problem would be necessary to avoid the possible emergence of a legal situation, where the tools we use today for privacy enhancing techniques are getting illegal.

6.2 Research Issues and Results

In this section we are going to describe the advancements of the research done by work package 10.0 in the fourth year of PRIME.

Performance of network layer anonymization. As a part of our work we studied the influence of several disturbing factors on the general performance of Tor circuits. Hence it was possible to find out limitations of nodes regarding various performance metrics, especially network latency and throughput.

We have compared the average latency of circuits created with the currently used method of path selection to uniformly chosen paths to measure the achieved improvements. Thus it is possible to justify the loss of anonymity introduced by the probabilistic path selection. Additionally, the influence of the length of paths on the performance of data transmissions was studied. Also we have shown that there is a medium positive correlation of circuit setup durations and the average measured latencies of the respective circuits.

Further, we have shown that by ensuring a low diversity of the routers in a path it is possible to lower the latencies of circuits, while one always has to keep in mind that the geographical diversity of the nodes in a path is an important substance to the security of the system. We therefore propose the integration of a geographical component into Tor clients that can be used in order to define lower and upper bounds of location diversity in paths, as well as any other geographical restrictions.

Additionally, we have proposed new methods that are based on active measurements of circuit latencies and passive throughput estimations in order to improve the overall performance of anonymous communication channels provided by Tor. Still, the latter need to be implemented and practically evaluated. Special emphasis has to be placed on possible implications of new methods on anonymity and security of the system.

Besides speeding up Tor, we also invested time into researching alternative anonymization techniques. First, we were looking into a method to provide an overlay network that does not need to hold states in its nodes of any kind. To this end, we proposed the usage of an UDP-based network layer together with ECC-encryption methods to create a suitable layer. As all hops of the packet require asymmetric encryption, we were specifically interested in reducing the amount of CPU-load on the nodes. The selected way was finally a novel trade-off between security and privacy that allowed clients as well as server nodes to both choose *independently* between anonymity and privacy on the client side and using RAM vs. using the CPU on the server side.

The limits of anonymization techniques. We have provided a calculation for the needed number of observations by a set of collaborating colluding network members in order to determine the rates of all honest members in Crowds anonymization system. Frequency consideration is important because it eliminates false positive categorization of peer partners, depending on the own threshold value for rate.

We have also shown that Crowds can not provide perfect security for its members under any admissible parameter values. The number of observations needed in order to determine the sending rate of the members in the crowd precisely enough can be relatively small. This rises a reasonable doubt on the possibility of using the system for strong anonymity in an open environment. Furthermore, we proposed an adaptive behaviour for the network participants to improve the provided degree of anonymity.

In another work, we showed that certain kinds of networks have additional information flow for the attacker, i.e. if they allow anonymous or pseudonymous replies to a user. This can be used to run all traditional attacks up to five times faster than in networks that do not have this additional feature.

Protection against the Sybil attack is a fundamental requirement for many of today's distributed applications. Online communities, peer-to-peer systems, and anonymous communication systems are only some well known examples. Providing privacy-friendly admission control for such applications is particularly challenging, especially without requiring online connectivity to a TTP. As a result of our work, we presented a solution to the Sybil attack that does not require online connectivity to a TTP *and* preserves user privacy. Our scheme is particularly suitable for being deployed in infrastructureless wireless networks. We have provided a security analysis of our proposal and discussed some real-world applications that would benefit from our solution. In particular, we showed how the anonymous communication mechanism Crowds can benefit from our approach when deployed in a mobile ad hoc network.

Holistic Anonymity. As the protection of an anonymizing network works only on the traffic layer, additional measures are necessary for data transferred on top, e.g. application layer data. To this end, credentials are a privacy enabled approach used to proof properties without revealing too much data. We could show however, that even given perfect security on the network layer and a strong credential system, an attacker can still learn enough information about users, possibly even enough to identify them.

An extension of our work on security of anonymizing networks that focused on traffic layer alone, will take into account information leakage on the application layer, too. Preliminary results show that if this kind of information is accessible in some form, attackers are not only able to build a profile of the user's interests at low costs, but also use this information to speed up traditional attacks that are targeted, e.g. at the user's peer partners.

Anonymization and Law. As research in the area of law enforcement in anonymization networks turned out to be practically impossible, we switched the view to research the impact of the data retention directive on traffic layer anonymization.

We found that the EC-directive differs from the actual implementations of the member states, and thus it is difficult to find a unified approach to this research topic. However, as only a couple of implementations were already in place, we could restrict our view to them. Findings [PK08] included that the directive was technically quite inaccurate, i.e. over-doing in some parts, while leaving other technical details disregarded.

6.3 Remaining Open Problems

As described before, many research issues and problems regarding anonymity providing communication infrastructure have been investigated and solved within the duration of the PRIME project, but some issues will need further investigation and new questions arise; some of which we will focus on in our future research. Besides researching basic questions related to security and performance, we will specifically try to answer questions that are closely related to anonymous communication. The latter refers to censorship resistant systems, appliance of network layer attacks on application layer, et al.

Practical Limits of Anonymous Communication. After having researched basic questions related to theoretical strength of anonymous communication systems, we will switch our focus to practical issues.

Due to the massive deployment of anonymizing networks in the recent time, it was made easier to check properties of these networks in the wild. Surprisingly, the most successful attacks were not targeted at the network layer, but either the application layer[For06], or side channel attacks[Mur06]. This shows dramatically that a lot of effort in traditional research was not targeted at the networks' weakest parts, while the attackers exploited exactly these.

Therefore, we will look for new holes and places of information leakage in these systems in order to fix them or find new solutions that avoid these. This includes, but is not limited to, information flow and user re-identification on the application layer, and secret ninja powers to get around the anonymous communication.

We will also like to look into how the average end users should behave in an real network in order to achieve a good degree of anonymity. While we know that turning off Java, Flash, and Javascript is necessary in order to avoid side channel attacks, no normal user will be willing to sacrifice a single bit of usability in order achieve some more privacy.

Extended usage of network layer attacks. Traditional attacks on anonymity networks are focused solely on the network layer. Yet, recent results have shown that due to their generic nature they are applicable to data transmitted on the application layer as well. Our estimation is that any attacks carried out on both layers will be more precise and much faster, thus requiring the users to take further actions in order to achieve a good level of security.

Performance of Anonymous Communication. While bandwidth seems to be available in abundance today, overlay networks whose concern is to retain the privacy of the users, typically suffer from the lack of bandwidth donations. This leads to a particularly high drop-out rate of connections and less bandwidth per user than in the normal Internet. Being confronted with the trade-off between ease of use and security, the typical end user chooses comfort and abandons usage of the anonymizing networks. To avoid this, we'll be looking for possibilities to guarantee a certain minimum amount of quality-of-service that would be acceptable by most privacy concerned users.

In order to achieve all that, we will try to adapt and validate known methods for routing, addressing and QoS for their suitability in anonymous peer-to-peer environments. Possibly, we will be able to build new approaches that better fit the faced needs, basing on the knowledge gained in the prior steps. In contrast to centralized solutions, which suffer from scalability problems, we will look into peer-to-peer based anonymity services to overcome this drawbacks. Furthermore, a decentralized approach could enable people like e.g. in China to have a possibility to use anonymization networks.

In parallel, mobile devices are getting more and more powerful to the extend, that they are more often used for data communication in addition to their original purpose, like e.g., placing telephone calls. On the other hand, their computational power, memory and esp. available bandwidth is most often insufficient to deploy solutions that have been proven successful in cabled networks. There are even additional risks, that are not given in the Internet, like information about person's location and the ease of eavesdropping wireless communication. In our research we are going to analyze those problems as additional constraints of the above topic.

In this area, we will closely collaborate with KU Leuven: first we will share anonymous traffic data needed to analyze the problem in depth. Together we will be looking for techniques that provide improved quality of service and quality of protection in anonymous networks. Thereafter we will discuss the results of our research with KU Leuven and commit findings about vulnerabilities and weaknesses in real networks as input for their work.

7 User-Side Identity Management Research (WP11.0)

The objective of PRIME and especially of work package 11.0 is to provide means and mechanisms for ensuring the users' right of self-determination with respect to processing their personal data. The fourth year of research in work package 11.0 has been a fruitful with respect to the five topics (1) technical quantification of privacy, (2) privacy-respecting reputation systems for future internet communities, (3) user-controlled privacy-enhancing identity management for multilateral interactions (4) externally provided information to support the user in managing identities, and (5) platform trust management and objective reputation management. All these research topics relate strongly to the project's aim of privacy-enhancing identity management.

7.1 Technical Quantification of Privacy

When a user discloses personal identifiable information (PII), her privacy decreases with respect to the service used. Often, it is not intuitively clear to a user how much the disclosure of certain PII affects her privacy. The goal of our research is to find measurements which can help the user in estimating the current privacy status. This estimation of the privacy status is meant to help the user in deciding what to do in situations where the user is given different options for disclosure of PII.

Within the last years, we have evaluated various approaches to quantifying anonymity and unlinkability which have been developed mainly focussing on communication layer for their applicability within a privacy-enhancing identity management system (PE-IMS). Based on that, we developed an approach for quantifying anonymity of users when performing actions supported by privacy-enhancing identity management. Further, we have developed an approach for quantifying linkability between different actions. Within the last project year, we have focussed on issues regarding efficiency of measurement and regarding usage of third parties for measurement. Further we have evaluated an alternative approach to estimating unlinkability properties using Formal Concept Analysis.

Research Results of the Fourth Project Year. After evaluating various approaches to quantifying anonymity and unlinkability known from literature and developing a quantification method directly addressing PE-IMS within the project years 1 to 3 (see [CCP07]), within the fourth project year we have focussed on two issues:

1. applicability and calculation efficiency of measurements for anonymity and unlinkability, and
2. an approach of estimating linkability properties using Formal Concept Analysis (FCA) [GW99]

Regarding item (1) we have analyzed complexity of measurement operations with respect to the number of attributes involved. Further, we have elaborated different ways to reduce complexity of calculations and storage, both with and without loss of correctness of measurement results. Also utilizing these efficiency improvements, we nevertheless need a large information base in order to get to useful measurement results for anonymity and unlinkability of user's actions.

Further, a major problem for calculation of anonymity metrics in the scope of PE-IMS is how to get enough information. While in principle this information is available from the users, it is privacy-relevant information, so users will not disclose it to a measuring entity except this entity (or these entities) can provide guarantees about privacy and security of the data.

So, in principle, there remain two options for a measuring entity: The first one is an approach distributing the calculation to different entities such that a subset of these entities does not get information about the sensitive data. The disadvantage of this approach is a further lack of efficiency. The second approach bases on tamper-resistant hardware and verified software, so that it can be verified that the measuring entity behaves as expected, and that this behaviour can be verified by third parties. The main problem here is verification of correctness and security of the hard- and software.

Summarizing this, we can say that the approach of using a trusted third party for privacy quantification is promising, but performance and trust issues need to be further evaluated within real world privacy-enhancing identity management systems in order to find out which approaches regarding performance, trust and security prove useful in practise. Research results regarding these issues have been published in [Cla07].

Regarding item (2) we have examined how conclusions about linkability threats can be drawn by analyzing message contents and subject knowledge in arbitrary communication systems. Therefore, we have defined messages and subjects as *formal contexts*. We have shown that concept lattices, which are achieved by applying Formal Concept Analysis to the concatenation of these formal contexts, can be used in order to draw conclusions about correlations, and therefore linkability, between contents of messages and knowledge of subjects.

This linkability analysis is particularly useful when the user of an identity management system is about to start a new action and needs support in choosing appropriate data items to disclose. With possible correlations to previous (actually unrelated) actions in mind, she can choose a reasonable set of data items, such that arising correlations cannot be used to re-identify her. This can even be supported in an automated way by means of a lattice structure.

An identity management system can moreover provide the capability to let the user browse through the concept lattice structure. That is, the user would be given the opportunity to explore her partial identities by means of concepts. The lattice structure provides, therefore, an easy to understand order of partial identities.

Besides development and evaluation of such a linkability estimation technique using FCA, which has been published in [BC07], we have also done reference implementation for this technique.

Remaining Open Problems. With regards to quantification of privacy within a PE-IMS several issues remain open for further research. After we now have models for application layer as well as for communication layer, more research needs to be done on an overall model combining both layers, especially the dependencies between the layers need to be modelled in more detail. The application layer model we developed does not take into account dependencies between different values of an attribute. So, more research can be done on taking advantage of such dependencies for privacy quantification. Further, changes of attribute values over time could be elaborated in more detail.

With regards to FCA-based modelling further research has to be done on the representation of all relevant correlations in single concepts. Suppose, two messages, for instance, one linkable by the adversary to a subject and the second message linkable (by a disjoint set of data items) to the first one. The link between this subject and the second message would not be derivable from a single concept. However, this affects just transitive linkability threats which arise from relations with one or more intermediate item.

7.2 Privacy-respecting Reputation Systems for Future Internet Communities

Internet communities gain more and more significance in the life of many people which use this convenient way to interact with people sharing the same interests. When interacting with strangers security requirements and trust issues become important.

Reputation systems have been designed and established in order to (1) help new interactors to estimate the others' behaviour and (2) motivate interactors to fulfil the others' expectations. Therefore, reputations systems collect and aggregate the experiences former interactors made in order to generate reputation metrics.

Unfortunately, besides generating trust between interactors the designs of reputation systems currently in use in Internet communities [Kol99] like eBay allow to generate user profiles including all contexts users have been involved in (e.g. time and frequency of participation, valuation of and interest in specific items). These profiles might become a promising target for numerous data collectors. However, this is contradictory to users' right of informational self-determination [MO04]. Although it has been rarely implemented so far it is possible to design reputation systems in a more privacy-respecting way.

The crucial point for the design of both a privacy-respecting and trustworthy reputation system is that users get at least partial control over the profiles built, and they are ensured that others cannot get rid of negative profiles.

To reach a compromise between different security requirements for reputation systems we developed the building blocks of pseudonym change with reputation transfer and parallel usage of pseudonyms/reputations depending on the context [Ste04, Ste06]. The building blocks were evaluated using the metrics for unlinkability of arbitrary items we presented in [SK03].

Research Results of the Fourth Project Year. Many people are members of various communities, typically even of communities with the same topics because of their interests. E.g., both eBay and Amazon are providers of Internet marketplace communities. For this reason users are interested in transferring reputation from one community to another community and using it there. Thus, there is the need for an interoperable reputation system that allows to transfer reputations between Internet communities. The privacy-respecting building blocks for reputation systems resulted in a system design that was implemented and evaluated by a diploma thesis [Pin07].

The building blocks also influenced the ENISA position paper on reputation systems [Ste] we contributed to.

There is still a need to enhance the first system design and discuss other design options of privacy-respecting and trustworthy reputation systems that allow interoperability between Internet communities. Especially an interdisciplinary discussion is needed.

7.3 User-Controlled Privacy-Enhancing Identity Management for Multilateral Interactions

User-controlled identity management enables trustworthy privacy protection when an individual interacts with an organization (customer-to-service, citizen-to-organization), or another individual (user-to-user). These are traditional bilateral scenarios. If interactions of any kind take place between arbitrary entities — individuals, organizations or subsets of them — we speak of multilateral interactions (MLI). In the scope of multilateral interaction scenarios, entities need to protect their privacy while being involved in various relations. Thus, extended requirements towards a privacy-enhancing identity management system evolve.

Research Results of the Fourth Project Year. During past research examining mechanisms which are available from traditional identity management systems, we identified building blocks for privacy-enhancing identity management for multilateral interactions (cf. [CCP07]). Those building blocks have been further developed in the course of the fourth project year and are outlined below according to their typical usage procedure:

1. **Pseudonyms and Partial identities:** Within the context of a user-controlled identity management system for multilateral interactions, pseudonyms and partial identities are used to control privacy of the individual entity. In accordance with the context, other involved parties, and objectives of an interaction, the entity should be able to specify recipients of its personal data and intended purpose or whether data should be revealed at all. Appropriate pseudonyms as identifiers for partial identities have to be chosen by the entity in order to control linkability of personal data disclosed.

2. **Relation Information:** Besides personal identifiable information belonging to exactly one entity, relation information exists that is not assignable to only one entity. Such information pertains to at least two entities and should be stored, managed, and evaluated in a privacy-sensitive way as well. In contrast to personal identifiable information where only intentions of one single owner have to be respected, relation information needs to take into account preferences from all parties involved.

3. **Selection of Interaction Partners:** Since collaboration and communication in a multilateral interaction environment represent direct interactions with other entities, the process of finding and selecting potential interaction partners is more complicated — particularly if it should be combined with mechanisms for user-controlled linkability. In general, the entity which is looking for collaboration and communication partners has two options: (1) it goes public, e.g. by self-advertising, or (2) it reverts to relations, e.g. the friend-of-a-friend principle.

4. **Trust Management and Reputation:** Reputation and trust may decisively influence the process of finding and selecting interaction partners in multilateral interactions. On the other hand, trust and reputation influence the negotiation of privacy policies between entities. In multilateral interactions, reputation does not have to be a static value assigned to an entity. It might also be of interest who has given the reputation and what kind of relation exists between this entity and the one that gets the reputation.

5. **Awareness:** Within the interaction environment awareness information exists which refers to

   a) the entity itself: its state, its activities and level of attention, etc. and thus awareness data may represent personal information.

   b) the interaction environment: available resources and functionalities, other entities, their availability and level of interest, etc.

   The identity management system should offer the entity possibilities to publish its own awareness information in a privacy-preserving way. Further, the system should allow access to awareness information of other entities in the environment in order to create or select an appropriate partial identity and to establish interactions.

6. **Context and History:** An entity in a multilateral interaction environment has varying goals towards different interaction partners at different times; we speak of a context of an interaction. In order to enhance privacy, the personal data of an entity is partitioned according to the respective context (cf. [FE06]). The selection of a partial identity will be made according to the current context and with consideration of history data, i.e., partial identities the entity has already used in the same or similar contexts and personal data which has already been disclosed in this scope.

7. **Access Control:** The use of access control lists or traditional role-based access control approaches is not possible within privacy-enhanced environments where users dynamically select partial identities in accordance with the current situation. Thus, a more flexible access control concept is required. We suggest an approach which is inspired by capabilities. However to avoid linkability of different partial identities of an entity, we use convertible credentials.

8. **Negotiation and Enforcement of Policies:** When entities interact with each other, different requirements for the handling of each partner's personal data arise. Depending on the other parties, preferences which data to disclose and which personal data from others to collect, execute, and store, may differ. Policies are negotiated according to the entity's own intentions with regard to other parties' preferences and legal regulations. Especially in multilateral interactions, negotiation processes may become complex since interests of several parties have to be considered. Besides the negotiation of privacy policies and preferences, mechanisms ensuring enforcement of the negotiated issues are required.

9. **Workflows and Behaviour Patterns:** Each entity possesses its own workflows and behaviour patterns which can be used to identify the entity and therefore represent personal identifiable information. In order to protect privacy, the identity management system should offer possibilities to avoid those individual workflows and behaviour patterns, for example, by adhering to templates. Templates are (1) either created beforehand by the system's designers or (2) generated by the entities themselves and offered to other entities not only for altruistic but egoistic reasons, namely to increase their privacy.

10. **External Guidelines and Regulations:** The privacy-enhancing identity management system should support the entity in managing, generating, and selecting appropriate partial identities. Therefore, not only data from the technical platform and other entities of the platform are necessary. Legal requirements and regulations play an important role as well. They have to be considered when decisions about the handling of one's own personal information or that of others are made. Since legislation differs between regions or countries and may change over time, this component always needs to be kept up-to-date.

The building blocks described above form a framework of components that need to be realized when developing a privacy-enhancing identity management system for multilateral interactions. However, ways of combining effective privacy protection with support for collaboration and cooperation while allowing meaningful relations between several entities at the same time still remain a research challenge and need to be investigated further.

Remaining Open Problems. A fundamental principle of user-controlled privacy emphasizes that each single entity has control over its personal information. This entity decides if, where, for what purpose, and how long its personal information is to be stored. Within the scope of multilateral interactions, the above hypothesis of privacy needs to be scrutinized and extended, since relation information is not only assigned to one but to two or more entities and represents personal information of all entities involved in that relation.

Our next steps for research in the field of multilateral interactions will focus on relation information and address questions such as

- How does relation information need to be handled?
- Will relation information only be stored if all involved parties agree?
- Where is relation information stored?
- Does the prohibition to store relation information affect user control of a single entity?
- Will related entities be asked or just informed if other partners disclose relation information?
- How could a negotiation process for storage and disclosure of relation information be designed and automated?
- How does relation information change over time?

This list only focuses on one building block. With regard to all others, more open issues and questions arise, that require further MLI research in general.

7.4 Externally Provided Information to Support the User in Managing Identities

As pure user-side solutions without co-operating parties are limited in identity management functionality, multilateral solutions are to be designed where the user gets information helping in successfully managing partial identities and their use. This information may be provided by transaction partners or by third parties (e.g., being integrated in the current transaction, or parties such as organizations or peers which enjoy certain confidence by the user). The provided information should support the user in privacy and identity management, in particular

1. getting an idea which knowledge on the own person other parties may have gained,
2. determining (re-)use of partial identities (taking into account e.g., the assumed or the stated trustworthiness of the service and its already compiled knowledge), and
3. deciding on future behaviour and configuration.

It should complement the picture which the Privacy-Enhancing Identity Management System can generate from past transactions ("Data Track" in the current PRIME Integrated Prototype) where transferred data and negotiated privacy policies have been logged. The research problem is the question how additional information provided by others can support the user and how specific information services providing this supporting information could work. In particular, the following research topics are addressed:

- Modelling of the integration of such information (e.g., which information, by which parties (organizations/peers), in which situation);

- Specifying formats and protocols to transmit this information to the PRIME software (e.g., by an RSS feed);
- Analysing the externally provided information with respect to relevance for the user;
- Communicating the information in an appropriate way to the user.

Research Results of the Fourth Project Year. In the first three years we analyzed and categorized information on privacy and security issues which can support the user in his identity management and implemented a security feed based on RSS for IPV1 (see the 2007 Research Report [CCP07]).

In the fourth year we put the security feed in the context of a variety of transparency tools which support the user, e.g., the "Data Track" as such or the possibility to exercize one's privacy rights online [FHPB+08, HFHPB07, Han08a, HCS08, Han08b]. Together with the team from the HCI work package some mock-ups were drafted.

In the last reports, it was already mentioned that security feeds can complement what Security Breach Notification Acts demand in the U.S. The importance of authentic security information will grow also in Europe as there are plans to enact a similar approach as part of an amended e-Privacy Directive, based on the study from Hogan & Hartson and Analysys 2006 [HA06] identifying a need of more transparency. This will foster approaches such as security feeds, but also other transparency tools.

The idea of the security feed – among other PRIME functionality – was contributed also to the work of the ENISA Working Group Privacy & Technology which will publish its deliverables in spring 2008.

7.5 Platform Trust Management and Objective Reputation Management

In this fourth period of PRIME, we have researched the use of privacy preferences as a means of user communicating to organizations how their personal information should be used and managed.

The key achievements have been two proof-of-concepts that have led to a greater understanding of how privacy preferences can be managed, and how they can support the objective-based Reputation Feedback Management that previous research had explored, and which forms a key component of the PRIME IPVx Platform Trust Management (PTM) module.

Objective Reputation Management. Previous research, reported more fully in the 2007 Research Report [CCP07], had examined how individuals could keep track of their information. They would be provided with tools that could examine the status of personal information that had previously been shared, either through an organization reporting compliance or through unsolicited requests made by the individual for status updates. This understanding allowed us to investigate the protocols that would enable such a range of interactions to take place, and on the basis of this understanding we have been able to build a full operational Proof of Concept (PoC). This PoC enabled the specific privacy concerns of a fictitious European agency that manages driver vehicle licensing to be explored.

The key problems that we wanted to address are:

1. What privacy preferences are most useful in to users?
2. How can privacy preferences be communicated to organizations is a hassle-free way?

3. How can evidence relating to compliance with privacy preferences be collated?

4. How should evidence of compliance be communicated to the users?

Based of the evidence gathers from the first PoC, we investigated at a more detailed level how individuals could manage their personal information and their privacy preferences, and the resulting *reputation* evidence. For our PoC we chose Microsoft's CardSpace Identity Meta Architecture (IMA) as the underlying messaging service. To the CardSpace Identity Selector we added a range of service to handle privacy preferences, both at the organization receiving the personal information (the Relying Party) and the Identity Provider which performs an identity brokering role. We also investigated cascaded Relying Parties and the added complexity that this created for the individual who is tracking their personal data. This resulted in new protocols for identity brokering.

We enhanced the Identity Selector, the client interface that enables individuals to choose which identity to share with the Relying Party, again taking the ID card metaphor that proved successful in our TrustGuide2 study [TC07]. The result is a fully operational identity management system incorporating privacy preferences. A final step was to provide Reputation Management for the client, which we did by developing the unique trust relationship that exists between the Identity Provider and the client. This enabled us to host our Reputation Manager at the Identity Provider, which could feedback objective reputation-building evidence to the client, using a standard web service/browser interface.

A full working prototype has been implemented, demonstrating the feasibility of our work and its applicability in complex enterprise environment. Some of the key findings of this recent research are:

1. Trust is central to engendering confidence and ensuring mass-market uptake of new technology.

2. Results have clearly confirmed our belief that giving individuals control over their PII engenders trust, but that achieving the right level(s) of controls is complex.

3. The ability to specify preferences at the time personal information is released is clearly a popular control, and one that can be easily understood by people if sensibly implemented. Preferences present a useful and pragmatic alternative to other privacy enhancing techniques, e.g. anonymity and pseudonymity.

4. Citizens are very adept at balancing convenience against risk.

5. Citizens demand clarity and transparency of use of their PII.

Papers and technical reports [TC07, CH07] have been published describing our results and findings. Future work in this area includes:

1. Publishing further papers and technical reports on the outcome of this research;

2. Further development around the selection of preferences including an advanced GUI.

3. Further research into the role that a trusted third party can play in helping users deal with the complexity of tracking and understanding how their personal user is being used.

8 Services-Side Identity Management Research (WP12.0)

8.1 Introduction

The goal of this work package is to research technologies and system solutions for automated privacy-respecting identity management in the services-side environment. The principal technical contributions towards privacy from the services-side are enhanced access control, obligation management, end-to-end privacy assurance control, personal data anonymization, private information retrieval, selective obfuscation of PII, and support for identity-aware devices. In addition, many of the issues in this space arise from the potentially highly-distributed nature of these "back-end" services and of the computing nodes on which they run. This raises challenges in the areas of service and platform credential and policy management and usage, of service and platform trustworthiness and of configuration compliance to requirements. This WP aims to advance the state of the art in all of these areas.

In this fourth period of PRIME, this WP addressed the following topics:

1. Scalable Obligation Management,
2. Identity-aware Devices,
3. End-to-End Assurance Control,
4. Person Data Anonymization,
5. Access Control for Privacy,
6. Private Information Retrieval and Anonymous Communication.

During this period, the principal achievements have been the further development of

- models and prototypes for scalable obligation management,
- a first working prototype of an identity-capable platform in a federated identity management context,
- models, protocols and prototypes for assurance control, for end-to-end user-driven checking of service-side capabilities for protection of user personal and sensitive data, both for B2C and B2B contexts,
- research in appropriate policy representation for assurance control, and more broadly for privacy policies that bridge the gap between user requirements and enforceable operational privacy policies,
- private information retrieval algorithm performance improvements,
- PRIME-compatible proof-of-concept prototypes for all of these.

The following next steps are required to progress those items:

1. large-scale trial deployments of the scalable obligation management system,
2. extension of identity-aware devices' functionalities, to handle the full lifecycle management of heterogeneous identity tokens, fully driven by policies and users' preferences,
3. development of a demonstrator making full use of access control enforcement, PII obfuscation and obligation mechanisms.

8.2 Scalable Obligation Management

The key problem addressed in this activity is how to manage obligation policies in a scalable way, on a potentially large set of personal data stored in various enterprise data repositories.

Previous HP Labs work consisted in researching and developing models and prototypes to handle privacy obligations, in organizational and enterprise contexts, consistently with current state-of-the-art identity management solutions and compatible with PRIME architecture and philosophy. Privacy obligations dictate duties and constraints on how to handle the lifecycle of personal data and digital identities (e.g. data deletion, notifications, data transformation, etc.), driven by laws, guidelines and users preferences. We soon understood that to make this happen, the obligation management system had to be scalable and perform in industrial environment.

Part of HP Labs' activity has been to understand and collate key requirements that need to be satisfied (based on customers' feedback, HP Labs analysis and lessons learnt):

1. Limit the number of "instantiated" obligation policies (and related management resources) independently on the amount of managed data and related privacy preferences;

2. Preserve the key capability to "customize" the management of each individual piece of personal data, based on users' privacy preferences;

3. Provide a more comprehensive automation of obligation policies, ensuring that obligations (once enforced) are not only passively monitored but also actions are taken to remediate/react to any violation. This to reduce the need for human intervention in case of large datasets.

Addressing this problem required researching on two key aspects:

1. how to represent "scalable" obligation policies;

2. how to manage, enforce and monitor these policies.

To address the stated problem and keep into account related requirements, we introduced the concept of parametric obligation policies [MB07c, MB07d]. A parametric obligation policy is a policy that leverages the concepts of our previous version of obligation policies but with the following key differences:

1. A parametric obligation policy can be associated to a potentially large set of personal data (i.e. no multiple instantiations) and, at the same time, it can dictate customized obligation constraints (based on users' privacy preferences) on each data item;

2. A parametric obligation policy does not embed privacy preferences in its Events and Actions sections (as instead happens in our previous version of obligation policies). Instead, this policy contains explicit references to these preferences, that are stored elsewhere — in data repositories;

3. The Target section [MB07c, MB07d] of parametric obligation policies explicitly model and describe the data repositories that will contain preference values pointed by these references — in addition to repositories containing personal data;

4. A new "On Violation" section [MB07c, MB07d] has been introduced to explicitly automate the process of "remediation" of violated obligations — as described in the requirement section.

The key feature introduced by parametric obligations is that privacy preferences are stored separately from parametric obligation policies: references are used to retrieve these preferences. This ensures that a parametric obligation policy can apply to a potentially large set of personal data — as defined in its Target element — and, at the same time, allows the "customization" of its Events and Actions based on references to external privacy preferences.

A model of scalable obligation policies and a related scalable obligation management system (SOMS) [MB07c, MB07d] has been researched and developed. The key innovation introduced in the SOMS system is its capability to dynamically interpret parametric obligation policies (i.e. their Target, Events, Actions and OnViolation Actions sections) and map their references on actual "targeted" data and preferences. This is done in an efficient way, via SQL queries that are instantiated on-the-fly — based on targeted data and related preferences.

A full working prototype has been implemented, demonstrating the feasibility of our work and its applicability in complex enterprise contexts.

Papers and technical reports [MB07c, MB07d] have been published describing our results and findings. Future work in this area includes:

1. publishing additional papers on this topic;
2. full integration of our work with state-of-the-art Enterprise Identity Management solutions (e.g. HP Identity Management suite);
3. further R&D work on graphical user interface to simplify administrative and obligation policy lifecycle management tasks.

8.3 Identity-aware Devices

The research problem addressed in this activity is how to leverage devices and federated identity management services to put users in control of their identities and enable a simpler, secure, trustworthy and transparent access to federated services and converged networks.

Current users' experience in networked and federated services is difficult and painful, especially when using mobile devices: users need to contact online service providers and authenticate against them, get credentials to access services and ensure these credentials are stored in a safe and secure place. Users have little control over the release their identity information and the related processes.

Specifically our work aims at addressing the following problems:

1. how to enable users to be in control of their identities and related "identity tokens" (e.g. identity credentials, identity attributes, access tokens/rights, etc.) by confidently and safely using (mobile) devices;
2. how to securely store, access and use "identity tokens" in "trustworthy" devices, driven by policies and user preferences, in a privacy compliant way;
3. how to enable a simpler, secure, trustworthy and transparent access to federated services and converged networks by using these devices.

Our research and development work in this space is consistent and aligned with recent PRIME work, that has focused on privacy management in the context of federated identity management.

A solution has been researched by HP Labs (jointly with HP business divisions and others) to address these issues, based on "Identity-aware Devices" (leveraging current devices, e.g. mobile phones, laptops, PDAs, etc.), trusted modules (driven by policies and users' preferences) and "Provisioning Services". This work has been carried out in collaboration with BT and Intel, in the context of a Liberty Alliance initiative, aiming at driving the next generation of interoperable identity solutions.

Solving these problems requires a solution that enables the vision of "Identity-aware Devices", where:

1. devices are safe and trustworthy "Personal Identity Providers", driven by policies;
2. devices can act on behalf of their users and/or other identity providers (for example, identity providers in a federated services scenario) via delegated "access tokens", and "certified identity tokens";
3. devices are provisioned with these tokens (and related policies) by trusted "Provisioning Services" that also manage their lifecycle;
4. users have a simplified, safe and secure interactions with federated services (including single-sign-on and identity management tasks), where communications with service and converged network providers are mediated by these "Identity-aware Devices", driven by policies and specified preferences.

Specifically, a working R&D solution has been developed, consisting of:

1. "Identity-aware Devices";
2. Federated (identity management) Services;
3. Registration and Provisioning Services, accessible as federated services.

The "Identity-aware Device" contains a "Trusted Module/Partition" (TM) hosting a Provisioned Module Manager (PMM) and one or more Provisioned Modules (PMs). The Provisioned Module Manager is a service running on the client platform which provides a "contact point" for provisioning operations, which can involve one or more Provisioned Modules.

The Provisioned Module is the key component which acts as the "Personal Identity Provider" and allows a simplified and controlled provisioning, management and release of identity tokens, based on policies and preferences. This is a trusted and secure component managing identities and identity tokens (manageable identities). This module has been designed to support:

1. Full lifecycle support for identity tokens: This includes support for transport medium independent (wired or over the air) provisioning, update, delete; activate, deactivate; serialize/de-serialize; and portability; etc;
2. Policy controlled access and operations: Which user can access which tokens; what can be done with each token; lifetime of a token; enforcement of user preferences (e.g. on black-listed service providers). Note that different types of "identity tokens" could be issued to (and managed by) the "Identity-aware Device", including: 802.1X wireless authentication tokens, VPN tokens, InfoCard/CardSpace tokens, OpenId tokens, SoftSim tokens, etc.

This Trusted Module (TM), embedded in devices, can be implemented either as a software or a hardware component. It can leverage Trusted Computing technology, such as TPM Modules, to provide: tamper resistant storage of tokens and sensitive data; Direct Anonymous Attestation (DAA) capabilities to engage in anonymous integrity checking and attestations of remote platforms.

In our current work, the “Identity-aware Device” supports the Liberty Alliance (LA) Identity Federated Framework (LA ID-FF) and its IDentity Web Service Framework (LA ID-WSF) standards, as well as the LA Active Client Technology (ACT) draft-standards. Federated Services are compliant to these LA standards.

The Registration Service is a federated, trusted service, compliant with the above LA standards, that allows a user to create a user account, express preferences and register one or more “identity-aware devices”. It also acts as a mediator to enable the provisioning of “identity tokens” (and related policies) to the device. The actual provisioning of these tokens is done by another trusted, federated service called Provisioning Service.

This Provisioning Service can be potentially co-located with the Identity Provider and is compliant with LA standards. This service ensures that the provisioning process can be carried out in two different phases: first phase consists in just issuing the “identity-aware device” with a “reference” to the identity token — as this identity token might not yet be required. In the next phase, “identity-aware device” engages with the Provisioning Service to de-reference this token. This enables flexibility and a convenient way to differentiate the token registration phase from the actual deployment/usage phase. Access to federated services happens via the mediation of the “identity-aware device”, subject to policy compliance and via LA federated protocols.

In this context, HP Labs have specifically focused on InfoCard/Cardspace tokens and explored how they can be provisioned within an identity-aware device and used by this device, by leveraging this framework.

A full working prototype has been developed previously and successfully demonstrated at RSA 2007. An additional demonstrator, based on this, has been developed by HP Labs to demonstrate the provisioning and use of InfoCard/Cardspace tokens, as an example of identity tokens.

This is work in progress: a second, enhanced demonstrator is in the process of being jointly developed and used in a technology trial. We are also exploring how this work can be further optimized using existing HP devices and security assets.

Future directions this research may evolve in could involve the following:

1. Further refinement of the overall management of policies within a identity-aware devices,

2. Further research how to fully factor in users’ preferences in identity-aware devices, by preserving its security and trustworthiness properties,

3. Further explore and research how to use other types of identity tokens within identity-aware devices.

8.4 Assurance Control

The research problem we consider is how to provide an automated policy compliance checking system which checks policies that are generated by the client. These client-generated policies may specify checking trust and assurance properties and judgments about trustworthiness of the service side, and they may take into account changing information at the IT resource level.

Solving this problem enables various usage cases, including:

- Users verifying whether the service side is capable of automatically executing privacy policies. This gives consumers the ability to determine whether unknown vendors on the Web are using IT systems and processes that can be trusted to execute their stated privacy policies.
- Automation of privacy assessment of the service side can be conveyed to the user in an open way (i.e., the compliance reports can be accessible to public) and with much more of a focus on evidence rather than having to rely on self-certification.

During this period, research by HP has produced:

1. A model of assurance control policies and a related assurance control system [EP07b, EP07a] has been researched and developed, and this is reflected within the PRIME architecture. The key innovation introduced in this system is to allow the decomposition of high-level privacy management requirements into automatically executable tests of detailed aspects of system configurations.
2. Proof-of-concept modules and sub-systems, including PRIME user-side and service-side modules for compliance checking [EP07b] that have been integrated into a B2B demonstrator and also within PRIME IPv3. These demonstrate the feasibility of our work and its applicability in enterprise contexts.
3. Refined UIs and associated user interactions for system policy compliance checking [FHPB$^+$08]; adaptations of these UIs have been used within the various implementations mentioned above; further details of the HCI research associated with this are given within the HCI Research section of this document.
4. Considering a broader range of scenarios, e.g., involving consideration of: a user-side SPCC [EP07b]; multiple services-side SPCCs [EP07a]; new protocols for compliance checking [EP07b]; protection during and after initial PII transfer using trusted computing [PMN08, PMN07, PMC07]; sticky policies [PM07]; links to Obligation Management [PM07, CMP07].
5. Research on developing a framework for bridging the gap between high-level and low-level policies using semi-formalized natural language policies to meet user requirements. The information shown in the user interface (UI) has to be understandable and actionable by a human user. Key research questions are how to best represent and decompose real life policies and legislation into actionable items and how to provide semi-automated ways of managing policies and controls that leverage existing policy language and enforcement frameworks. Our solution centres around the definition of an intermediate layer of 'privacy-positive' policies in human readable form. One type of such policies are the assurance control user-side and service-side policies we define; another are policies that represent privacy requirements and activities and controls which can be carried out to meet these requirements.

This solution can be exploited by integrating it with state-of-the-art enterprise identity management and/or audit solutions, to enhance existing functionality in order to allow a greater degree of system policy compliance checking. During 2007, this technology has been integrated into the PRIME identity manager, and HP Labs has also demonstrated its feasibility by integrating this technology with an enterprise audit solution, so that low-level information about system properties and behaviour can be collected and fed into the system, which then is able to assess the level of satisfaction of higher-level privacy and security statements, and separately with user-side and service-side obligation management systems. Papers and technical reports [EP07b, EP07a] have been published describing our results and findings.

Future work in this area may include:

1. publishing additional papers on this topic,
2. integration of our work with state-of-the-art Enterprise Identity Management and audit solutions,
3. further R&D work on graphical UIs for assurance control,
4. further work on bridging the gap between user requirements and enforceable operational privacy policies.

8.5 Personal Data Anonymization

In the healthcare domain, as in many other domains, it is useful to analyze and process personal data (such as patient records), e.g., for statistical or scientific research purposes, without disclosing the identity of the related persons. On this topics, previous PRIME research has dealt with the capability to control the linkability between data issued by different organizations (spatial linkability) and/or at different times (temporal linkability) for different kinds of end users, but also on the capability for the patient to control the creation of anonymized personal data, and to control the reversibility of the anonymization process, by using a personal smartcard. A demonstrator has been included in PRIME prototype V0 by LAAS, and an article on this topic has been submitted in 2007 as a chapter for a book on Privacy Technologies for Healthcare [KD07].

8.6 Access Control for Privacy

Our research on this topic deals with access control enforcement of privacy policies. LAAS developed an access control scheme where a request for a composite operation (i.e., an operation involving one or several accesses to PII) must first be authorized by the Access Control Decision Function (ACDF). If the composite operation is authorized by the ACDF, authorization proofs are generated for all the elementary accesses that can be part of the composite operation. These authorization proofs are then checked by an access control interceptor at each access into PII. The results of this research have been integrated in the PRIME architecture, and are the base of the development of PRIME Integrated Prototype access control. This year, the main research efforts were devoted to the generation of authorization proofs for particular composite operations.

Other protection techniques are currently analyzed to complement access control mechanisms, in particular to protect PII data against possible abuse by privileged users (e.g., security administrator, system operator, etc.) or against misuse by faulty software. We analyzed and compared techniques such as sticky policies and identity-based cryptography developed by HP Labs, policy-based encryption developed by Eurecom, and fragmentation-scattering techniques previously developed by LAAS. We are implementing a derived IPV3 prototype where some of these techniques will be experimented, as well as other aspects that have not been experimented in PRIME yet, including PII obfuscation.

More oriented to controlling accesses between servers belonging to different organizations, we have made a survey on security models and policies on Access Control for Collaborative Systems [BKD07].

8.7 Private Information Retrieval and Anonymous Communications

Private Information Retrieval (PIR) is a cryptographic technique used to protect privacy of requests to a database, including against a malicious or at least curious database manager. In this context, the fact that a given user is interested in a given information is by itself private. For example, requests to a public web server providing information on particular diseases such as AIDS can be considered sensitive. If the web server is using PIR, any observer, even a malicious administrator of the web server, is unable to identify which data are queried and retrieved by the user.

A survey of current PIR algorithms and a comparison of their performance has enabled us to identify that the PIR performance analyzes proposed by different authors were not suited for comparing these algorithms. This has led us to propose a new metric for PIR algorithms, which is generic enough to cope with all peculiarities of each PIR algorithm. This metrics and a comparison between these algorithms have been published in 2006, altogether with an analysis on how some PIR algorithms can be extended to be applied on statistical databases.

We have also analyzed how to use PIR algorithms to develop a new kind of mix (that we call pMIX), which is more efficient than conventional ones. In particular, the PIR technique guarantees that a malicious mix administrator or any external observer is unable to relate incoming and outgoing data flows. Thus, by using this technique, it is possible to provide anonymous communications through a single pMIX rather than through a network or a cascade of mixes. Diverse variants of this solution have been developed and their performance has been analyzed. We have shown that it is possible to obtain anonymous communications with very low latency, bringing anonymity to new interactive Internet applications, such as Voice on IP, which were incompatible with the high latency induced by conventional mixes [AMD07, AMDIC07].

Currently, we are trying to design new PIR algorithms with significantly reduced overhead for the applications we consider. One of these algorithms has been published [AMG07], which reduces the PIR computation cost by two orders of magnitude with respect to previously published algorithms. We expect more improvements to achieve about three orders of magnitude.

A book on Ambient Intelligence has been published in 2007, with a chapter on Security, Privacy and Availability [DPR07]. Moreover, several keynote talks on Privacy have been given in 2007 by Yves Deswarte, including one at ICISS 2006 in Delhi and one at the French STIC conference [Des07].

9 Conclusion

In the fourth year of the PRIME project, the six research work packages have made significant progress towards understanding, modelling and developing the techniques and components necessary for a privacy-friendly identity management system.

Towards this goal, WP6.1 has worked towards easy to understand user interfaces supporting users in using privacy-enhancing technologies with special focus on suitable choices of predefined preference settings providing the most privacy-friendly options and protecting the privacy interests of most end users. WP8.0 has developed different methods for providing location privacy with regards to location based services and selective data protection methods for distributed data storage. WP9.0 has continued to develop new cryptographic tools and definitions related to identity management in general and anonymity in particular, WP10.0 has improved the quality of service of anonymizing networks, analyzed limits of its security, and has done research regarding holistic anonymity and law enforcement in anonymization networks. WP11.0 has continued to analyze and specify privacy metrics for identity management systems, third party services proving privacy-relevant information to support identity selection, as well as privacy aspects regarding multilateral interactions and reputation systems. WP12.0 has refined tools for server-side scalable obligation management, identity-aware devices and assurance control and has done research regarding privacy-compliant access control and private information retrieval.

The reported results appear in numerous publications, including leading conferences and journals. We have included abstracts for these published results in the appendix.

As outlined in each of the sections, there are still issues open to further research beyond the PRIME project.

These include improvement of easy-to-use yet informative user interface design; improved methods for privacy of location information; efficient fragmentation and encryption solutions for remote data storage; new cryptographic mechanisms with regards to delegation of credentials, searchable encryption, and oblivious service access; enhancing security and performance of holistic anonymization mechanisms; improvements in user-side and services-side identity management methods, specific transparency techniques for end users, reputation systems, assurance control and scalability analysis for the system components.

Some of these challenges will be taken up by one of the follow up projects of PRIME, PrimeLife. PrimeLife will be 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.

References

[ACD$^+$07a] Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, and Pierangela Samarati. Location privacy protection through obfuscation-based techniques. In *Proc. of the 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security*, Redondo Beach, CA, USA, July 2007.

[ACD$^+$07b] Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, and Pierangela Samarati. Managing privacy in LBAC systems. In *Proc. of the Second IEEE International Symposium on Pervasive Computing and Ad Hoc Communications (PCAC-07)*, Niagara Falls, Canada, May 2007.

[ACD$^+$07c] Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, and Pierangela Samarati. A middleware architecture for integrating privacy preferences and location accuracy. In *Proc. of the 22nd IFIP TC-11 International Information Security Conference (SEC 2007)*, Sandton, South Africa, May 2007.

[ACD$^+$07d] Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, and Pierangela Samarati. Privacy-enhanced location services information. In A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, and C. Lambrinoudakis, editors, *Digital Privacy: Theory, Technologies and Practices*. Taylor and Francis Group, 2007.

[ACDS07a] Claudio A. Ardagna, Marco Cremonini, Sabrina De Capitani di Vimercati, and Pierangela Samarati. Location privacy in pervasive computing. In S. Gritzalis, T. Karygiannis, and C. Skianis, editors, *Security and Privacy in Mobile and Wireless Networking*. Troubador Publishing, 2007.

[ACDS07b] Claudio A. Ardagna, Marco Cremonini, Sabrina De Capitani di Vimercati, and Pierangela Samarati. Privacy-enhanced location-based access control. In M. Gertz and S. Jajodia, editors, *The Handbook of Database Security: Applications and Trends*. Springer-Verlag, 2007.

[ACDS07c] Claudio A. Ardagna, Marco Cremonini, Sabrina De Capitani di Vimercati, and Pierangela Samarati. Privacy in the electronic society: Emerging problems and solutions. In *ISI Platinum Jubilee Monograph Series*. World Scientific Press, 2007.

[AGGS08] André Adelsbach, Ulrich Greveler, Stephan Groß, and Sandra Steinbrecher. ANOCAST: Rethinking broadcast anonymity in the case of wireless communications. In *Proceedings of Sicherheit 2008*, GI Lecture Notes in Informatics, 2008.

[AKMP08] Christer Andersson, Markulf Kohlweiss, Leonardo Martucci, and Andriy Panchenko. Self-certified and sybil-free framework for secure digital identity domain buildup. In *Proceedings of the Workshop in Information Security Theory and Practices 2008: Smart Devices, Convergence and Next Generation Networks (WISTP 2008)*, Sevilla, Spain, May 2008.

[AMD07] Carlos Aguilar-Melchor and Yves Deswarte. Anonymous Voice over IP communications. Technical Report 07664, LAAS-CNRS, November 2007.

[AMDIC07] Carlos Aguilar-Melchor, Yves Deswarte, and Julien Igutchi-Cartigny. Closed-circuit unobservable Voice over IP. In *23rd Annual Computer Security Applications Conference — ACSAC 2007*, pages 119–128, Miami Beach, FL, USA, December 2007. IEEE-CS. Available at http://www.acsac.org/2007/papers/110.pdf.

[AMG07] Carlos Aguilar-Melchor and Philippe Gaborit. A lattice-based computationally-efficient private information retrieval protocol. In *Western European Workshop on Research in Cryptology — WEWoRC 2007*, Bochum, Germany, July 2007. An extended version is available at http://eprint.iacr.org/2007/446.pdf.

[AMSS07] Antonia Azzini, Stefania Marrara, Roberto Sassi, and Fabio Scotti. A multimodal fuzzy approach for biometric authentication. In *Proc. of the International Conference on Knowledge-Based and Intelligent Information & Engineering Systems*, Vietri sul Mare, Italy, September 2007.

[AP07] Christer Andersson and Andriy Panchenko. Practical anonymous communication on the mobile internet using tor. In *Proceedings of the Third International Workshop on the Value of Security through Collaboration (IEEE SECOVAL 2007) part of IEEE SECURECOMM'07*, Nice, France, September 2007.

[BC07] Stefan Berthold and Sebastian Clauß. Linkability estimation between subjects and message contents using formal concepts. In Atsuhiro Goto, editor, *DIM '07, Proceedings of the 2007 ACM Workshop on Digital Identity Management*, pages 36–45, November 2007.

[BCKL08] Mira Belenkiy, Melissa Chase, Markulf Kohlweiss, and Anna Lysyanskaya. P-signatures and noninteractive anonymous credentials. In *Theory of Cryptography Conference (TCC 2008)*, Lecture Notes in Computer Science, page 18, New York,NY,USA, 2008. Springer-Verlag.

[Ber08] Mike Bergmann. Generic predefined privacy preferences for online applications. In Simone Fischer Hübner, Penny Duquenoy, Albin Zuccato, and Leonardo Martucci, editors, *The Future of Identity in the Information Society: Proceedings of the Third IFIP WG 9.2, 9.6/11.6, 11.7/FIDIS International Summer School*, International Summerschool Karlstad, Sweden, May 2008. Springer.

[BFK00] Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable internet access. In H. Federrath, editor, *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*, pages 115–129. Springer-Verlag, LNCS 2009, July 2000.

[BKD07] Amine Baina, Anas Abou El Kalam, and Yves Deswarte. Access control for collaborative systems: a state of the art survey. Technical Report 07119, LAAS-CNRS, March 2007.

[BMBS07] Adrian Baldwin, Marco Casassa Mont, Yolanta Beres, and Simon Shiu. On identity assurance in the presence of federated identity management systems. In *Proceedings of DIM'07 - 2007 ACM Workshop on Digital Identity Management*, pages 27–35. ACM, 2007.

[BMS07] Adrian Baldwin, Marco Casassa Mont, and Simon Shiu. On identity assurance in the presence of federated identity management systems. In *HPL-2007-47*, HPL Technical Reports. HP, 2007.

[BPHL$^+$07] Katrin Borcea-Pfitzmann, Marit Hansen, Katja Liesebach, Andreas Pfitzmann, and Sandra Steinbrecher. Managing ones identities in organisational and social settings. Datenschutz und Datensicherheit, Vieweg-Verlag 31/9, S. 671-675., 2007.

[BPS00] Oliver Berthold, Andreas Pfitzmann, and Ronny Standtke. The disadvantages of free MIX routes and how to overcome them. In H. Federrath, editor, *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*, pages 30–45. Springer-Verlag, LNCS 2009, July 2000.

[BS02] A. Beimel and Y. Stahl. Robust information-theoretic private information retrieval. In S. Cimato C. Galdi G. Persiano, editor, *Proceedings of the 3rd Conference on Security in Communication Networks*, volume 2576 of *Lecture Notes in Computer Science*, pages 326–341. Springer-Verlag, 2002.

[CAG02] Lorrie F. Cranor, M. Arjula, and P. Guduru. Use of a P3P user agent by early adopters. *In Proceedings of the ACM Workshop on Privacy in the Electronic Society*, November 2002.

[CCP07] Jan Camenisch, Sebastian Clauß, and Andreas Pfitzmann, editors. *PRIME Research Report.* PRIME, 2007.

[CDF$^+$07] Valentina Ciriani, Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati. Fragmentation and encryption to enforce privacy in data storage. In *Proc. of the 12th European Symposium On Research In Computer Security*, Dresden, Germany, September 2007.

[CDFS07] Valentina Ciriani, Sabrina De Capitani di Vimercati, Sara Foresti, and Pierangela Samarati. k-anonymous data mining: A survey. In Charu C. Aggarwal and Philip S. Yu, editors, *Privacy-Preserving Data Mining: Models and Algorithms.* Springer-Verlag, 2007.

[CG05] Lorrie F. Cranor and Simon Garfinkel. *Security and Usability: Designing Secure Systems that People Can Use.* O'Reilly, 2005.

[CGKS95] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. In *FOCS '95: Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS'95)*, page 41, Washington, DC, USA, 1995. IEEE Computer Society.

[CGP$^+$07] Stelvio Cimato, Marco Gamassi, Vincenzo Piuri, Roberto Sassi, and Fabio Scotti. A biometric verification system addressing privacy concerns. In *Proc. of the IEEE*

International Conference on Computational Intelligence and Security, Harbin, China, December 2007.

[CH07] Elisabetta Carrara and Giles Hogben, editors. *Reputation-based Systems: a security analysis.* ENISA, 2007. ENISA Position Paper No. 2.

[CHP07] Jan Camenisch, Susan Hohenberger, and Michael Østergaard Pedersen. Batch verification of short signatures. In Moni Naor, editor, *EUROCRYPT*, volume 4515 of *Lecture Notes in Computer Science*, pages 246–263. Springer, 2007.

[Cla07] Sebastian Clauß. *Towards Quantification of Privacy Within a Privacy-Enhancing Identity Management System.* PhD thesis, Technische Universität Dresden, December 2007.

[CMP07] Stephen Crane, Marco Cassasa Mont, and Siani Pearson. On helping individuals to manage privacy and trust. In *Trust Management.* Icfai University Press, June 2007.

[CNas07] Jan Camenisch, Gregory Neven, and abhi shelat. Simulatable adaptive oblivious transfer. In Moni Naor, editor, *Advances in Cryptology - EUROCRYPT 2007*, volume 4515 of *Lecture Notes in Computer Science*, pages 573–590, Barcelona,E, 2007. Springer-Verlag.

[Dan03] George Danezis. Statistical disclosure attacks: Traffic confirmation in open environments. In Gritzalis, De Capitani di Vimercati, Samarati, and Katsikas, editors, *Proceedings of Security and Privacy in the Age of Uncertainty, (SEC2003)*, pages 421–426, Athens, May 2003. IFIP TC11, Kluwer.

[DD07] George Danezis and Claudia Diaz. Space-efficient private search. In Rachna Dhamija and Sven Dietrich, editors, *Proceedings of Financial Cryptography (FC2007)*, volume 4886 of *Lecture Notes in Computer Science*, pages 148–162, Tobago, 2007. Springer-Verlag.

[DD08] George Danezis and Claudia Diaz. A survey of anonymous communication channels. Microsoft technical report, Microsoft Research, 2008.

[DDF$^+$07a] Ernesto Damiani, Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati. An experimental evaluation of multi-key strategies for data outsourcing. In *Proc. of the 22nd IFIP TC-11 International Information Security Conference (SEC 2007)*, Sandton, South Africa, May 2007.

[DDF$^+$07b] George Danezis, Claudia Diaz, Sebastian Faust, Emilia Ksper, Carmela Troncoso, and Bart Preneel. Efficient negative databases from cryptographic hash functions. In Juan A. Garay, Arjen K. Lenstra, and Masahiro Mambo, editors, *Proceedings of the 10th Information Security Conference (ISC 2007)*, volume 4779 of *Lecture Notes in Computer Science*, pages 423–436, Valparaiso,Chile, 2007. Springer-Verlag.

[DDT07] George Danezis, Claudia Diaz, and Carmela Troncoso. Two-sided statistical disclosure attack. In Nikita Borisov and Philippe Golle, editors, *Proceedings of Privacy Enhancing Technologies, 7th International Workshop, PET 2007*, volume

4776 of *Lecture Notes in Computer Science*, pages 30–44, Ottawa,Canada, 2007. Springer-Verlag.

[Des07] Yves Deswarte. Intelligence ambiante et protection de la vie privée. In *Grand Colloque STIC 2007*, Paris, France, November 2007. Available at http://www.rntl.org/colloqueSTIC2007/docs/Presentations/IntAmb\&PVP-Y-Dewarte.pdf.

[DFJ+07a] Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati. A data outsourcing architecture combining cryptography and access control. In *Proc. of the 1st Computer Security Architecture Workshop*, Fairfax, Virginia, USA, November 2007.

[DFJ+07b] Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, and Pierangela 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.

[DFJ+07c] Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati. Privacy of outsourced data. In A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, and C. Lambrinoudakis, editors, *Digital Privacy: Theory, Technologies and Practices*. Taylor and Francis Group, 2007.

[DFJS07] Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, and Pierangela Samarati. Access control policies and languages. *International Journal of Computational Science and Engineering*, 3(2):94–102, 2007.

[DFPS07] Sabrina De Capitani di Vimercati, Sara Foresti, Stefano Paraboschi, and Pierangela Samarati. Access control models for XML. In M. Gertz and S. Jajodia, editors, *The Handbook of Database Security: Applications and Trends*. Springer-Verlag, 2007.

[DJPS07] Sabrina De Capitani di Vimercati, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati. Trust management services in relational databases. In *Proc. of the ACM Symposium on InformAtion, Computer and Communications Security (ASIACCS'07)*, Singapore, March 2007.

[DMS04] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, 2004.

[DP08] Mina Deng and Bart Preneel. On secure and anonymous buyer-seller watermarking protocol. In *International Conference on Internet and Web Applications and Services*, page 8, Athens,Greece, 2008.

[DPR07] Yves Deswarte, David Powell, and Yves Roudier. *Sécurité, protection de la vie privée et disponibilité*, chapter 13, pages 301–344. Arago 31. Hermès-Lavoisier, 2007. Information at http://www.lavoisier.fr/notice/gb/not2.asp?id=36ONXOZ3SRLOFJ.

[DS07] Sabrina De Capitani di Vimercati and Pierangela Samarati. Data privacy: Problems and solutions. In *Proc. of the Third International Conference on Information Systems Security (ICISS 2007)*, Delhi, India, December 2007.

[DS08] George Danezis and Len Sassaman. How to bypass two anonymity revocation schemes. In Nikita Borisov and Ian Goldberg, editors, *Privacy Enhancing Technologies Symposium*, Lecture Notes in Computer Science, page 15, Leuven,BE, 2008. Springer-Verlag.

[DTD07] Claudia Diaz, Carmela Troncoso, and George Danezis. Does additional information always reduce anonymity? In Ting Yu, editor, *Workshop on Privacy in the Electronic Society 2007*, pages 72–75, Alexandria,VA,USA, 2007. ACM.

[EP07a] Tariq Ehsan Elahi and Siani Pearson. Privacy assurance: Bridging the gap between preference and practice. In Costas Lambrinoudakis, Günther Pernul, and A Min Tjoa, editors, *Trust, Privacy and Security in Digital Business*, volume 4657 of *LNCS*, pages 65–74. Springer-Verlag, 2007.

[EP07b] Tariq Ehsan Elahi and Siani Pearson. Towards enhanced privacy assurance mechanisms and systems for on-line participants. In *HPL-2006-189*, HPL External Technical Reports. HP, January 2007.

[FCP07] Alberto Ferrante, Satish Chandra, and Vincenzo Piuri. A query unit for the IPSec databases. In *Proc. of SECRYPT-2007*, Barcelona, Spain, July 2007.

[FE06] Elke Franz and Benjamin Engel. A realization of context management facilitating the usage of partial identities. In *Proceedings of PEP, CHI Workshop on Privacy-Enhanced Personalization 2006*, 2006.

[FGSB08] Elke Franz, Christin Groba, Thomas Springer, and Mike Bergmann. A comprehensive approach for context-dependent privacy management. In *Proccedings of the Third International Conference on Availability, Reliability and Security - ARES 4-7 March 2008*, Polytechnic University of Catalonia, Barcelona, Spain, 2008. IEEE Computer Society.

[FHP08] Simone Fischer-Hübner and John Sören Pettersson. Usable privacy and identity management: Challenges and approaches. In *Proceedings of the IADIS International Conference ICT, Society and Human Beings 22 - 24 July 2008*, Polytechnic University of Catalonia, Barcelona, Spain, 2008.

[FHPB$^+$08] Simone Fischer-Hübner, John Sören Pettersson, Mike Bergmann, Marit Hansen, Siani Pearson, and Marco Casassa Mont. Chapter 11: HCI designs for privacy-enhancing identity management. In Alessandro Acquisti, Stefanos Gritzalis, Costos Lambrinoudakis, and Sabrina De Capitani Di Vimercati, editors, *Digital Privacy: Theory, Technologies, and Practices*, pages 229–252. Auerbach Publications, Taylor & Francis Group, 2008.

[FJK06] S.M. Furnell, A. Jusoh, and D. Katsabas. The challenges of understanding and using security: A survey of end-users. *Computers & Security*, Vol. 25(1):27–35, 2006.

[For06] FortConsult. Practical onion hacking: Finding the real address of tor clients, October 2006.

[FP07] Alberto Ferrante and Vincenzo Piuri. High-level architecture of an IPSec-dedicated system on chip. In *Proc. of the 3rd EuroNGI Conference on Next Generation Internet Networks NGI 2007*, Trondheim, Norway, May 2007.

[Fur04a] S.M. Furnell. E-commerce security: a question of trust. *Computer Fraud & Security*, pages 10–14, October 2004.

[Fur04b] S.M. Furnell. Using security: easier said than done? *Computer Fraud & Security*, pages 6–16, April 2004.

[Fur05] S.M. Furnell. Why users cannot use security. *Computers& Security*, Vol. 24(4):274–279, 2005.

[Ger02] Daniela Gerd tom Markotten. User-centered security engineering. *Proceedings of the 4th EurOpen/USENIX Conference NordU2002*, February 2002. http://www.iig.uni-freiburg.de/telematik/atus/publications.html.

[Gol07] Ian Goldberg. Improving the robustness of private information retrieval. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, May 2007.

[GS05] Oliver Günther and Sarah Spiekermann. RFID and the perception of control: The consumer's view. *Communications of the ACM*, 48(9):73–76, September 2005.

[GS08] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel Smart, editor, *EUROCRYPT*, 2008. To appear.

[GW99] B. Ganter and R. Wille. *Formal Concept Analysis: Mathematical Foundations*. Springer-Verlag, Berlin, Heidelberg, 1999.

[HA06] Hogan & Hartson and Analysys. Preparing the next steps in regulation of electronic communications – a contribution to the review of the electronic communications regulatory framework, 2006.

[Han08a] Marit Hansen. Marrying transparency tools with user-controlled identity management. In Simone Fischer-Hübner, Penny Duquenoy, Albin Zuccato, and Leonardo Martucci, editors, *The Future of Identity in the Information Society: Proceedings of the Third IFIP WG 9.2, 9.6/11.6, 11.7/FIDIS International Summer School*, International Summerschool Karlstad, Sweden, May 2008. Springer.

[Han08b] Marit Hansen. User-controlled identity management: key to the future of privacy? *International Journal of Intellectual Property Management (IJIPM), Special Issue on: "Identity, Privacy and New Technologies"*, 2008. to appear.

[HCS08] Marit Hansen, Alissa Cooper, and Ari Schwartz. Privacy and identity management. *IEEE Security & Privacy*, May/June 2008. to appear.

[HFHPB07] Marit Hansen, Simone Fischer-Hübner, John Sören Pettersson, and Mike Bergmann. Transparency tools for user-controlled identity management. In *Expanding the Knowledge Economy: Issues, Applications, Case Studies; Proceedings of eChallenges 2007*, pages 1360–1367, Amsterdam, 2007. IOS Press.

[JG00] Uwe Jendricke and Daniela Gerd tom Markotten. Usability meets security — the identity-manager as your personal security assistant for the internet. *Proceedings of the 16th Annual Computer Security Applications Conference*, December 2000. http://tserv.iig.uni-freiburg.de/telematik/forschung/projekte/kom_technik/atus/publications.html.

[KD07] Anas Abou El Kalam and Yves Deswarte. Secure anonymization of healthcare records. Technical Report 07426, LAAS-CNRS, August 2007. Submitted as a chapter for the book ”Assimilating Privacy Technologies and Heath Care Compliance”, Publisher: IDEA Group Inc.

[KFF$^+$07] Markulf Kohlweiss, Sebastian Faust, Lothar Fritsch, Bartek Gedrojc, and Bart Preneel. Efficient oblivious augmented maps: Location-based services with a payment broker. In Nikita Borisov and Philippe Golle, editors, *Proceedings of Privacy Enhancing Technologies, 7th International Workshop, PET 2007*, volume 4776 of *Lecture Notes in Computer Science*, page 20, Ottawa,Canada, 2007. Springer-Verlag.

[Kol99] Peter Kollock. The production of trust in online markets. *Advances in Group Processes*, 16, 1999.

[KPP07] Dogan Kesdogan, Vinh Pham, and Lexi Pimenidis. Information disclosure in identity management. In *Proceedings of 12th Nordic Workshop on IT-Security, NordSec*, Reykjavik, Iceland, Oct 2007.

[KPP08] Dogan Kesdogan, Vinh Pham, and Lexi Pimenidis. Analyse der Verkettbarkeit in nutzergesteuertem Identitätsmanagement. In *Proceedings of Sicherheit 2008, 4. Jahrestagung, Fachbereich Sicherheit der Gesellschaft für Informatik*, Saarbruecken, Germany, April 2008.

[KPS07] Klaus Kursawe, Peter Palfrader, and Len Sassaman. Echolot and Leuchtfeuer: Measuring the reliability of unreliable mixes. Technical Report ESAT-COSIC 2007-005, ESAT-COSIC, 2007.

[KZSD07] Eleni Kosta, Jan Zibuschka, Tobias Scherner, and Jos Dumortier. Privacy-enhancing user-friendly identity management for location based services using PRIME technology - a legal discussion. In *Cyberlaw Security & Private, The Second International Conference on Legal, Security and Privacy Issues in Information Technology (LSPI)*, Beijing, China, 2007.

[LPW$^+$07] Olaf Landsiedel, Lexi Pimenidis, Klaus Wehrle, Heiko Niedermayer, and Georg Carle. More: A peer-to-peer based connectionless onion router. In *Proceedings of IEEE GLOBECOM, Globalcommunications Conference*, Washington DC, USA, November 26th 2007.

[Mat04] Nick Mathewson. Pynchon Gate Protocol draft specification, September 2004. http://www.abditum.com/pynchon/.

[MB07a] Marco Casassa Mont and Boris Balacheff. On device-based identity management in enterprises. In *HPL-2007-53*, HPL Technical Reports. HP, 2007.

[MB07b] Marco Casassa Mont and Boris Balacheff. On device-based identity management in enterprises. In Costas Lambrinoudakis, Gunter Pernul, and A. Min Tjoa, editors, *Trust, Privacy and Security in Digital Businesses*, volume 4657 of *LNCS*, pages 94–103. Springer, 2007.

[MB07c] Marco Casassa Mont and Filipe Beato. On parametric obligation policies: Enabling privacy-aware information lifecycle management in enterprises. In *Proceedings of the 8th IEEE International Workshop on Policies for Distributed Systems and Networks 2007*, pages 51–55. IEEE, 2007.

[MB07d] Marco Casassa Mont and Filipe Beato. On parametric obligation policies: Enabling privacy-aware information lifecycle management in enterprises. In *HPL-2007-7*, HPL Technical Reports. HP, 2007.

[MD04] Nick Mathewson and Roger Dingledine. Practical traffic analysis: Extending and resisting statistical disclosure. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, volume 3424 of *LNCS*, pages 17–34, May 2004.

[MKAP08] Leonardo Martucci, Markulf Kohlweiss, Christer Andersson, and Andriy Panchenko. Self-certified Sybil-free Pseudonyms: Introducing privacy in infrastructureless wireless networks. In *Proceedings of 1st ACM Conference on Wireless Network Security (WiSec 2008)*, Alexandria, USA, 31 Mar – 2 Apr 2008.

[MO04] Tobias Mahler and Thomas Olsen. Reputation systems and data protection law. In *eAdoption and the Knowledge Economy : Issues, Applications, Case Studies*, pages 180–187, Amsterdam, 2004. IOS Press.

[Mon07] Marco Casassa Mont. Automation of privacy management. *Digma Magazine*, March 2007.

[Mur06] Steven J. Murdoch. Hot or not: Revealing hidden services by their clock skew. In *13th ACM Conference on Computer and Communications Security (CCS)*, Alexandria, Virginia, USA, 2006.

[Nie04] Jacob Nielsen. Jacob Nielsen's Alertbox, user education is not the answer to security problems, October 2004. http://www.useit.com.

[OS05] R. Ostrovsky and W.E. Skeith, III. Private searching on streaming data. In *CRYPTO 2005*, pages 223–240, 2005.

[Pet08] John Sören Pettersson. Little brother is watching you — commercialisation of personal data through 'webification'. In *Proceedings of the colloquium of project Asphalès, Paris November 22-23, 2007: "Security of the digitized man" (Réflexions prospectives et internationales: "La sécurité de l'individu numérisé"*, Paris, French, 2008.

[PFH07] John Sören Pettersson and Simone Fischer-Hübner. *Human Values of IT*, chapter Transparency as the key to user-controlled personal data processing. Cambridge Scholars Press, 2007.

[PFHND^{+}05] John Sören Petterson, Simone Fischer-Hübner, Jenny Nilsson Ninni Danielsson, Mike Bergmann, Thomas Kriegelstein, Sebastian Clauss, and Henry Krasemann. Making PRIME usable. In *SOUPS*, Carnegie Mellon University, USA, July 2005. ACM Digital Library.

[Pim07] Lexi Pimenidis. On anonymity loss in open environments. In *Proceedings of 12th Nordic Workshop on IT-Security, NordSec*, Reykjavik, Iceland, Oct 2007.

[Pin07] Franziska Pingel. Realisierung eines datenschutzfreundlichen Reputationssystems mit Anschluss an die technische Infrastruktur phpBB für Internet Communities. Diplomarbeit, Institut für Systemarchitektur, Technische Universität Dresden, 2007.

[PJS$^+$08] Andreas Pfitzmann, Andreas Juschka, Anne-Katrin Stange, Sandra Steinbrecher, and Stefan Köpsell. Communication privacy. In Alessandro Acquisti, Stefanos Gritzalis, Costos Lambrinoudakis, and Sabrina De Capitani Di Vimercati, editors, *Digital Privacy: Theory, Technologies, and Practices*, pages 19–46. Auerbach Publications, Taylor & Francis Group, 2008.

[PK08] Lexi Pimenidis and Eleni Kosta. The impact of the retention of traffic and location data on the internet user. *Datenschutz und Datensicherheit, Zeitschrift für rechts- und prüfungssicheres Datenmanagement*, 32(02):92 – 97, 2008.

[PM07] Siani Pearson and Marco Casassa Mont. A system for privacy-aware resource allocation and data processing in dynamic environments. In *HPL-2006-185*, HPL Technical Reports. HP, 2007.

[PMC07] Siani Pearson, Marco Cassasa Mont, and Stephen Crane. Analysis of trust properties and related impact of trusted platforms. In *Trust Management*. Icfai University Press, June 2007.

[PMN07] Siani Pearson, Marco Casassa Mont, and Manny Novoa. Securing information transfer within distributed computing environments. In *HPL-2007-10*, HPL Technical Reports. HP, 2007.

[PMN08] Siani Pearson, Marco Casassa Mont, and Manny Novoa. Securing information transfer within distributed computing environments. *IEEE Security & Privacy Magazine*, Jan/Feb 2008.

[PP07a] Andriy Panchenko and Lexi Pimenidis. Crowds revisited: Practically effective predecessor attack. In *Proceedings of the 12th Nordic Workshop on Secure IT-Systems (NordSec 2007)*, Reykjavik, Iceland, October 2007.

[PP07b] Andriy Panchenko and Lexi Pimenidis. Using trust to resist censorship in the presence of collusion. In *Proceedings of 22nd IFIP TC-11 International Information Security Conference*, Sandton Convention Centre, Sandton, South Africa, May 2007.

[PP07c] David R. Piegdon and Lexi Pimenidis. Targeting physically addressable memory. In *Fourth GI International Conference on Detection of Intrusions & Malware, and Vulnerability Assessment*, pages 193–212, 2007.

[PPR08] Andriy Panchenko, Lexi Pimenidis, and Johannes Renner. Performance analysis of anonymous communication channels provided by tor. In *Proceedings of the Third International Conference on Availability, Reliability and Security (ARES 2008)*, Barcelona, Spain, March 2008.

[PS07] Meredith L. Patterson and Len Sassaman. Subliminal channels in the private information retrieval protocols. In *Proceedings of the 28th Symposium on Information Theory in the Benelux*, page 8, Enschede,NL, 2007. Werkgemeenschap voor Informatie- en Communicatietheorie.

[PSC08] Meredith L. Patterson, Len Sassaman, and David Chaum. Freezing more than bits: Chilling effects of the OLPC XO security model. In *Proceedings of the Workshop on Usability, Psychology, and Security 2008*. USENIX, 2008.

[Sas07] Len Sassaman. The faithless endpoint: How tor puts certain users at greater risk. Technical report esat-cosic 2007-003, ESAT-COSIC, 2007.

[SCM05] Len Sassaman, Bram Cohen, and Nick Mathewson. The Pynchon Gate: A secure method of pseudonymous mail retrieval. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2005)*, Arlington, VA, USA, November 2005.

[Sco07] Fabio Scotti. Computational intelligence techniques for reflections identification in iris biometric images. In *Proc. of the IEEE International Conference on Computational Intelligence for Measurement Systems and Applications*, Ostuni, Italy, June 2007.

[Sim83] Gustavus J. Simmons. The prisoners' problem and the subliminal channel. In *CRYPTO*, pages 51–67, 1983.

[SK03] Sandra Steinbrecher and Stefan Köpsell. Modelling unlinkability. In Roger Dingledine, editor, *Proceedings of Privacy Enhancing Technologies workshop (PET 2003)*, pages 32–47. Springer-Verlag, LNCS 2760, 2003.

[SP07a] Len Sassaman and Bart Preneel. The byzantine postman problem: A trivial attack against PIR-based nym servers. Technical Report ESAT-COSIC 2007-001, ESAT-COSIC, 2007.

[SP07b] Len Sassaman and Bart Preneel. Solving the byzantine postman problem. Technical Report ESAT-COSIC 2007-004, ESAT-COSIC, 2007.

[Ste] Sandra Steinbrecher. Privacy-respecting reputation systems for future internet communities. The European e-Identity conference, ENISA Workshop on Security Issues in Reputation Systems, 12-13 June 2007, Paris, France.

[Ste04] Sandra Steinbrecher. Balancing privacy and trust in electronic marketplaces. In *Proceedings of Trust and Privacy in Digital Business, First International Conference, TrustBus*, volume 3184 of *Lecture Notes in Computer Science*, pages 70–79. Springer, 2004.

[Ste06] Sandra Steinbrecher. Design options for privacy-respecting reputation systems within centralised Internet communities. In *Proceedings of IFIP Sec 2006, 21st IFIP International Information Security Conference: Security and Privacy in Dynamic Environments*, May 2006.

[TC07] Ellen Tweney and Stephen Crane. Trustguide2: An exploration of privacy preferences in an online world. In Paul Cunningham, editor, *eChallenges 2007*, volume Volume 4, Part 2 of *Information and Communication Technologies and the Knowledge Economy*, pages 1379–1385. IOS Press, 2007.

[TDDP07] Carmela Troncoso, Claudia Diaz, Orr Dunkelman, and Bart Preneel. Traffic analysis attacks on a continuously-observable steganographic file. In T. Furon et al., editor, *Information Hiding, 9th International Workshop*, volume 4567 of *Lecture Notes in Computer Science*, pages 220–236, Saint-Malo,FR, 2007. Springer-Verlag.

[Wat05] Brent Waters. Efficient identity-based encryption without random oracles. In *EUROCRYPT '05*, volume 3494 of LNCS, pages 320–329, 2005.

[WDH07] Brecht Wyseur, Mina Deng, and Thomas Herlea. A survey of homomorphic encryption schemes. COSIC internal report, COSIC, 2007.

[WT99] A. Whitten and J.D. Tygar. Why Jonny can't encrypt: A usability evaluation of PGP 5.0. In *Proceedings of the Ninth USENIX Security Symposium*, 1999.

[ZFR^{+}07a] Jan Zibuschka, Lothar Fritsch, Mike Radmacher, Tobias Scherner, and Kai Rannenberg. Enabling privacy of real-life LBS: A platform for flexible mobile service provisioning. In *in Proc. of the 22nd IFIP TC-11 International Information Security Conference (SEC 2007)*, Sandton, South Africa, May 2007.

[ZFR^{+}07b] Jan Zibuschka, Lothar Fritsch, Mike Radmacher, Tobias Scherner, and Kai Rannenberg. Privacy-friendly LBS: A prototype-supported case study. In *Proc. of the 13th Americas Conference on Information Systems*, Keystone, Colorado, USA, August 2007.

[ZRSR07] Jan Zibuschka, Mike Radmacher, Tobias Scherner, and Kai Rannenberg. Empowering LBS users: Technical, legal and economic aspects. In *Proc. of the eChallenges conference*, The Hague, The Netherlands, October 2007.

A Collection of Abstracts of Publications

This section provides the abstracts of all the manuscripts written by the work packages WP6.1 and WP[8-12].0 during the reporting period. Most of them have already been published, some are to appear or in submission to conferences and journals. Some publications are joint work between different WPs, which is indicated with the respective publications. The abstracts of these publications are only provided once.

A.1 WP6.1

1. Simone Fischer-Hübner, John Sören Pettersson, Mike Bergmann, Marit Hansen, Siani Pearson, Marco Casassa Mont: *HCI Designs for Privacy-Enhancing Identity Management* [FHPB+08].
   Chapter 11 in *Digital Privacy: Theory, Technologies, and Practices.*
   (Joint work with WP11.0 and WP12.0.)
   Abstract. Privacy-enhancing Identity Management can enable users to retain and maintain informational self-determination in our networked society. This chapter reports results from the HCI (Human Computer Interaction) research work that has been done within the European Union FP6 project on Privacy and Identity Management for Europe (PRIME) within its first three project years. It discusses and compares alternative user interface (UI) paradigms for privacy-enhancing Identity Management, and shows how legal privacy principles derived from the European Union Directives as well as social requirements for enhancing trust by end users have been mapped into suggestions of user interface solutions.

2. John Sören Pettersson, Simone Fischer-Hübner, *Transparency as the key to user-controlled personal data processing* [PFH07].
   In: *Human Values of IT (to appear).*
   Abstract. A fundamental legal privacy principle is transparency: a user must be able to see who is processing his personal data and for what purpose. With today's Internet technologies it should not be difficult for people to monitor their data at different databases. There are furthermore laws stating under what conditions people can revoke their data. However, in usability tests we have noted that people are unaware of their rights and are in general uncertain if the technology can give reliable trust indicators. The chapter discusses these problems for making people confident information society citizens, and presents tools for enhancing transparancy and making it possible for people to exercise their legal rights. We also note that the some European laws, for instance, the Swedish Personal Data Act, do not account for citizens being computer users; the laws for exercising the right to access one's own data free of charge is sometimes granted only 'once per annum'. This limitation is justified in a society where information exchange is carried out by paper-based mail, but is questionable when automated services are omnipresent and people have many and varied encounters with services every year.

3. Mike Bergmann, *Generic Predefined Privacy Preferences for Online Applications* [Ber08].
   In: *The Future of Identity in the Information Society: Proceedings of the Third IFIP WG 9.2, 9.6/11.6, 11.7/FIDIS International Summer School.*
   Abstract. Every day users disclose various kinds of personal data using the Internet for daily activities. The disclosed data in summary may draw a perfect picture of them. Up to now it is difficult for end users to decide what to disclose and what to hide.We try to

support the user in this task and propose a limited set of applicable predefined privacy preferences taking privacy principles into account.We will apply these preferences for typical online activities to evaluate and to enhance them. We elaborate the dependencies and correlations between the privacy preferences and application scenarios. As a final result and based on the proposed privacy preferences we introduce a privacy-enhancing data disclosure splitting guiding the user step by step through the process of data disclosure.

4. Marit Hansen, *Marrying Transparency Tools with User-Controlled Identity Management* [Han08a].
   In: *"The Future of Identity in the Information Society": Proceedings of the Third IFIP WG 9.2, 9.6/11.6, 11.7/FIDIS International Summer School.*
   (Joint work with WP11.0, please refer to Section A.5 for the abstract.)

5. Marit Hansen, Simone Fischer-Hübner, John Sören Pettersson, Mike Bergmann, *Transparency Tools for User-Controlled Identity Management* [HFHPB07].
   In: *Proceedings of the eChallenges 2007 conference.*
   (Joint work with WP11.0.)
   Abstract. Transparency is an important precondition for the users' control over their privacy. It can increase the users' trust in accurate and secure processing of their personal data. This paper presents concepts and implementations of different transparency tools, which are employed in a user-controlled identity management system of the project PRIME – Privacy and Identity Management for Europe. A focus is laid on showing the concepts and their visualization via user interfaces. Considerations on human-computer interaction are highlighted for different alternatives, and the motivation for design choices is explained. Moreover, results from user tests are reported and analyzed. In the following section related work is described. This paper concludes that currently transparency tools rarely belong to standard functionality of information and communication technologies although much improvement would already be achievable today.

6. John Sören Pettersson, *"Little Brother Is Watching You- commercialisation of personal data through webification"* [Pet08].
   In: *"Security of the digitized man" (Réflexions prospectives et internationales: "La sécurité de l'individu numérisé") the concluding colloquium of project Asphalés.*
   To appear in proceedings published in 2008 by CNRS with an abstract in French.

7. Elke Franz, Christin Groba, Thomas Springer, Mike Bergmann, *"A Comprehensive Approach for Context-dependent Privacy Management"* [FGSB08].
   In: *Proccedings of the Third International Conference on Availability, Reliability and Security - ARES 2008.*
   Abstract. Disclosing personal and contextual information enables a widespread use of services in the internet in a comfortable working environment adapted to individual user needs and preferences. On the other hand, information disclosure raises serious privacy issues. In this paper, we propose a novel integrated approach for achieving a good balance between convenience and privacy. We integrate Privacy- Enhancing Identity Management (PIM) with extended rolebased access control mechanisms (RBAC) in an extended concept and architecture. In combination, both mechanism enable a fine grained control of what private and contextual information is disclosed to whom in what situation. We have implemented a Privacy-Enhanced Adaptive Communicator application (PEAC) to validate our approach.

8. Simone Fischer-Hübner, John Sören Pettersson, *"Usable Privacy and Identity Management: Challenges and Approaches"* [FHP08].
In: *Proceedings of the IADIS International Conference ICT (invited keynote presnetation), Society and Human Beings 2008.*
Abstract. A critical success factor for Privacy-Enhancing Technologies (PETs) will be user-friendly and intelligible user interfaces that convey and enhance trust. Such user interfaces have to meet challenges such as the user-friendly representation of complex PET concepts, such as "pseudonyms", "unlinkabilty" or "anonymous credentials" that are unfamiliar to many users, the provision of security against phishing or spoofing attacks, the enforcement of legal privacy principles, such as informed consent or transparency, as well as the mediation of reliable trust to the end users. In this presentation, we will in discuss such challenges for usable and privacy-enhancing identity management and will provide some HCI guidelines for addressing those challenges.

A.2 WP8.0

1. Sabrina De Capitani di Vimercati and Sara Foresti and Sushil Jajodia and Pierangela Samarati: *Access Control Policies and Languages* [DFJS07].
In *International Journal of Computational Science and Engineering.*
Abstract. Access control is the process of mediating every request to data and services maintained by a system and determining whether the request should be granted or denied. Expressiveness and flexibility are top requirements for an access control system together with, and usually in conflict with, simplicity and efficiency. In this paper, we discuss the main desiderata for access control systems and illustrate the main characteristics of access control solutions.

2. Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, Pierangela Samarati: *A Data Outsourcing Architecture Combining Cryptography and Access Control* [DFJ+07a].
In *Proc. of the 1st Computer Security Architecture Workshop.*
Abstract. Data outsourcing is becoming today a successful solution that allows users and organizations to exploit external servers for the distribution of resources. Some of the most challenging issues in such a scenario are the enforcement of authorization policies and the support of policy updates. Since a common approach for protecting the outsourced data consists in encrypting the data themselves, a promising approach for solving these issues is based on the combination of access control with cryptography. This idea is in itself not new, but the problem of applying it in an outsourced architecture introduces several challenges.

   In this paper, we first illustrate the basic principles on which an architecture for combining access control and cryptography can be built. We then illustrate an approach for enforcing authorization policies and supporting dynamic authorizations, allowing policy changes and data updates at a limited cost in terms of bandwidth and computational power.

3. Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, Pierangela Samarati: *Over-encryption: Management of Access Control Evolution on Outsourced Data* [DFJ+07b].
In *Proc. of the 33rd International Conference on Very Large Data Bases (VLDB 2007).*
Abstract. Data outsourcing is emerging today as a successful paradigm allowing users

and organizations to exploit external services for the distribution of resources. A crucial problem to be addressed in this context concerns the enforcement of selective authorization policies and the support of policy updates in dynamic scenarios.

In this paper, we present a novel solution to the enforcement of access control and the management of its evolution. Our proposal is based on the application of selective encryption as a means to enforce authorizations. Two layers of encryption are imposed on data: the inner layer is imposed by the owner for providing initial protection, the outer layer is imposed by the server to reflect policy modifications. The combination of the two layers provides an efficient and robust solution. The paper presents a model, an algorithm for the management of the two layers, and an analysis to identify and therefore counteract possible information exposure risks.

4. Valentina Ciriani, Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, Pierangela Samarati: *Fragmentation and Encryption to Enforce Privacy in Data Storage* [CDF+07].
   In *Proc. of the 12th European Symposium On Research In Computer Security.*
   Abstract. Privacy requirements have an increasing impact on the realization of modern applications. Technical considerations and many significant commercial and legal regulations demand today that privacy guarantees be provided whenever sensitive information is stored, processed, or communicated to external parties. It is therefore crucial to design solutions able to respond to this demand with a clear integration strategy for existing applications and a consideration of the performance impact of the protection measures.

   In this paper we address this problem and propose a solution to enforce privacy over data collections by combining data fragmentation with encryption. The idea behind our approach is to use encryption as an underlying (conveniently available) measure for making data unintelligible, while exploiting fragmentation as a way to break sensitive associations between information.

5. Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, Pierangela Samarati: *Location Privacy Protection Through Obfuscation-based Techniques* [ACD+07a].
   In *Proc. of the 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security.*
   Abstract. The widespread adoption of mobile communication devices combined with technical improvements of location technologies are fostering the development of a new wave of applications that manage physical positions of individuals to offer location-based services for business, social or informational purposes. As an effect of such innovative services, however, privacy concerns are increasing, calling for more sophisticated solutions for providing users with different and manageable levels of privacy. In this work, we propose a way to express users privacy preferences on location information in a straightforward and intuitive way. Then, based on such location privacy preferences, we discuss a new solution, based on obfuscation techniques, which permits us to achieve, and quantitatively estimate through a metric, different degrees of location privacy.

6. Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, Pierangela Samarati: *Managing Privacy in LBAC Systems* [ACD+07b].
   In *Proc. of the Second IEEE International Symposium on Pervasive Computing and Ad Hoc Communications (PCAC-07).*

Abstract. One of the main challenges for privacy-aware location-based systems is to strike a balance between privacy preferences set by users and location accuracy needed by Location-Based Services (LBSs). To this end, two key requirements must be satisfied: the availability of techniques providing for different degrees of user location privacy and the possibility of quantifying such privacy degrees. To address the first requirement, we describe two obfuscation techniques. For the second requirement, we introduce the notion of relevance as the estimator for the degree of location obfuscation. This way, location obfuscation can be adjusted to comply with both user preferences and LBS accuracy requirements.

7. Ernesto Damiani, Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, Pierangela Samarati: *An Experimental Evaluation of Multi-key Strategies for Data Outsourcing* [DDF+07a].
In *Proc. of the 22nd IFIP TC-11 International Information Security Conference (SEC 2007).*
Abstract. Data outsourcing is emerging today as a successful solution for organizations looking for a cost-effective way to make their data available for on-line querying. To protect outsourced data from unauthorized accesses, even from the (honest but curious) host server, data are encrypted and indexes associated with them enable the server to execute queries without the need of accessing cleartext. Current solutions consider the whole database as encrypted with a *single key* known only to the data owner, which therefore has to be kept involved in the query execution process. In this paper, we propose different *multi-key data encryption* strategies for enforcing access privileges. Our strategies exploit different keys, which are distributed to the users, corresponding to the different authorizations. We then present some experiments evaluating the quality of the proposed strategies with respect to the amount of cryptographic information to be produced and maintained.

8. Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, Pierangela Samarati: *A Middleware Architecture for Integrating Privacy Preferences and Location Accuracy* [ACD+07c].
In *Proc. of the 22nd IFIP TC-11 International Information Security Conference (SEC 2007).*
Abstract. Location-Based Access Control (LBAC) systems support the evaluation of conditions on locations in the enforcement of access control policies. The ability to evaluate conditions on a set of authorized locations has a number of well-known advantages, including enriching access control expressiveness. However, when locations are used in combination with personal identities, users privacy must be considered. In this paper, we describe a solution to integrate a LBAC system with privacy-enhanced techniques based on location obfuscation. Our solution is based on a privacy-aware middleware component that explicitly addresses the trade-off between users privacy and location accuracy by satisfying preferences set by users and maximizing the quality of location information released to LBAC systems.

9. Sabrina De Capitani di Vimercati, Sushil Jajodia, Stefano Paraboschi, Pierangela Samarati: *Trust Management Services in Relational Databases* [DJPS07].
In *Proc. of the ACM Symposium on InformAtion, Computer and Communications Security (ASIACCS'07).*
Abstract. Trust management represents today a promising approach for supporting ac-

cess control in open environments. While several approaches have been proposed for trust management and significant steps have been made in this direction, a major obstacle that still exists in the realization of the benefits of this paradigm is represented by the lack of adequate support in the DBMS.

In this paper, we present a design that can be used to implement trust management within current relational DBMSs. We propose a trust model with a SQL syntax and illustrate the main issues arising in the implementation of the model in a relational DBMS. Specific attention is paid to the efficient verification of a delegation path for certificates. This effort permits a relatively inexpensive realization of the services of an advanced trust management model within current relational DBMSs.

10. Valentina Ciriani, Sabrina De Capitani di Vimercati, Sara Foresti, Pierangela Samarati: *k-Anonymous Data Mining: A Survey* [CDFS07].
In *Privacy-Preserving Data Mining: Models and Algorithms.*
Abstract. Data mining technology has attracted significant interest as a means of identifying patterns and trends from large collections of data. It is however evident that the collection and analysis of data that include personal information may violate the privacy of the individuals to whom information refers. Privacy protection in data mining is then becoming a crucial issue that has captured the attention of many researchers.

    In this chapter, we first describe the concept of k-anonymity and illustrate different approaches for its enforcement. We then discuss how the privacy requirements characterized by k-anonymity can be violated in data mining and introduce possible approaches to ensure the satisfaction of k-anonymity in data mining.

11. Sabrina De Capitani di Vimercati, Sara Foresti, Pierangela Samarati: *Access Control Models for XML* [DFPS07].
In *The Handbook of Database Security: Applications and Trends.*
Abstract. XML has become a crucial tool for data storage and exchange. In this chapter, after a brief introduction on the basic structure of XML, we illustrate the most important characteristics of access control models. We then discuss two models for XML documents, pointing out their main characteristics. We finally present other proposals, describing their main features and their innovation compared to the previous two models.

12. Claudio A. Ardagna, Marco Cremonini, Sabrina De Capitani di Vimercati, Pierangela Samarati: *Privacy-enhanced Location-based Access Control* [ACDS07b].
In *The Handbook of Database Security: Applications and Trends.*
Abstract. Advancements in location technologies reliability and precision are fostering the development of location-based services that make use of the location information of users. An increasingly important category of such services is represented by Location-based Access Control (LBAC) systems that integrate traditional access control mechanisms with access conditions based on the physical position of users and other attributes related to the users location. Since privacy is extremely important for users, protection of their location information is paramount to the success of such emerging location-based services.

    In this chapter, we first present an overview of Location-based Access Control systems and then characterize the location privacy protection problem. We then discuss the main techniques that have been proposed to protect location information, focusing on the obfuscation-based techniques. We conclude the chapter by showing a privacy-aware

LBAC architecture and describing how a location-based access control policy can be evaluated.

13. Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, Pierangela Samarati: *Privacy in the Electronic Society: Emerging Problems and Solutions* [ACDS07c].
In *ISI Platinum Jubilee Monograph Series.*
Abstract. As the global information infrastructure is becoming more ubiquitous, digital business transactions are increasingly performed using a variety of mobile devices and across multiple communication channels. This new service-oriented paradigm is making the protection of privacy an increasing concern, as it relies on rich context representations (e.g., of location and purpose) and requires users to provide a vast amount of information about themselves and their behavior. This information is likely to be protected by a privacy policy, but restrictions to be enforced may come from different input requirements, possibly under the control of different authorities. In addition, users retain little control over their personal information once it has been disclosed to third parties. Secondary usage regulations are therefore increasingly demanding attention. In this paper, we present the emerging trends in the data protection field to address the new needs and desiderata of today's systems.

14. Sabrina De Capitani di Vimercati, Pierangela Samarati: *Data Privacy: Problems and Solutions* [DS07].
In *Third International Conference on Information Systems Security (ICISS 2007).*
Abstract. Nowadays, the global information infrastructure connects remote parties worldwide through the use of large scale networks, relying on application level protocols and services such as the World Wide Web. The vast amounts of personal information thus available has led to growing concerns about the privacy of their users. In this paper, we briefly discuss some privacy issues that have to be considered to address the new needs and desiderata of today's systems and discuss ongoing work.

15. Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, Pierangela Samarati: *Privacy of Outsourced Data* [DFJ+07c].
In *Digital Privacy: Theory, Technologies and Practices.*
Abstract. None.

16. Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, Pierangela Samarati: *Privacy-enhanced Location Services Information* [ACD+07d].
In *Digital Privacy: Theory, Technologies and Practices.*
Abstract. None.

17. C.A. Ardagna, M. Cremonini, S. De Capitani di Vimercati, P. Samarati: *Location Privacy in Pervasive Computing* [ACDS07a].
In *Security and Privacy in Mobile and Wireless Networking.*
Abstract. None.

18. Alberto Ferrante and Vincenzo Piuri: *High-level Architecture of an IPSec-dedicated System on Chip* [FP07].
In *Proc. 3rd EuroNGI Conference on Next Generation Internet Networks NGI 2007.*
Abstract. IPSec is a suite of protocols which adds security to communications at the IP level. Protocols within the IPSec suite make extensive use of cryptographic algorithms.

Since these algorithms are computationally very intensive, some hardware acceleration is needed to support high throughput. In this paper we propose a high level architecture of a System on Chip (SoC) which implements IPSec. This SoC has been thought to be placed on the main data path of the host machine (ow-through architecture), thus allowing for transparent processing of IPSec trafc. The functionalities of the different blocks and their interactions, along with an estimation of the internal memory size, are also shown.

19. Alberto Ferrante, Satish Chandra, Vincenzo Piuri: *A Query Unit for the IPSec Databases* [FCP07].
In *SECRYPT 2007*.
Abstract. IPSec is a suite of protocols that adds security to communications at the IP level. Protocols within IPSec make extensive use of two databases, namely the Security Policy Database (SPD) an d the Security Association Database (SAD). The ability to query the SPD quickly is fundamental as this operation needs to be done for each incoming or outgoing IP packet, even if no IPSec processing needs to be applied on it. This may easily result in millions of query per second in gigabit networks. Since the databases may be of several thousands of records on large secure gateways, a dedicated hardware solution is needed to support high throughput. In this paper we discuss an architecture for these query units, we propose different query methods for the two databases, and we compare them through simulation. Two different versions of the architecture are presented: the basic version is modified to support multithreading. As shown by the simulations, this technique is very effective in this case. The architecture that supports multithreading allows for 11 million queries per second in the best case.

20. Antonia Azzini, Stefania Marrara, Roberto Sassi, Fabio Scotti: *A Multimodal Fuzzy Approach for Biometric authentication* [AMSS07].
In *Proc. of the International Conference on Knowledge-Based and Intelligent Information & Engineering Systems.*
Abstract. In the last few years the security of the user's identity has become of paramount importance. In this paper we investigate the opportunity of using a multimodal biometric system as input of a fuzzy controller designed with the aim of preventing user substitution after the initial authentication process.

21. Fabio Scotti: *Computational intelligence techniques for reflections identification in iris biometric images* [Sco07].
In *Proc. of the IEEE International Conference on Computational Intelligence for Measurement Systems and Applications.*
Abstract. Iris biometric systems identify individuals by comparing the characteristics of the iris acquired by the acquisition sensors. When reflections are present in the iris image, the portion of the image covered by reflections must be discarded from any further comparison since it can produce false matches. The paper presents a methodology for reflections identification in iris biometric images based on neural networks. In particular, this paper proposes a set of features which can be extracted from the iris image and that can be effectively used to achieve an accurate identification of the reflection position using a neural network. Moreover, the paper presents how the radial symmetry operator can be used as a proper feature to identify the reflections in iris images. The method is general and can be used in any biometric system based on iris images.

22. Stelvio Cimato, Marco Gamassi, Vincenzo Piuri, Roberto Sassi, Fabio Scotti: *A biometric*

verification system addressing privacy concerns [CGP+07].
In *Proc. of the IEEE International Conference on Computational Intelligence and Security.*
Abstract. Biometric techniques are more and more exploited in order to fasten and make more reliable the identification process. Recently, many proposals have been formulated combining cryptography and biometrics in order to increase the confidence in the system when biometric templates are stored for verification. In this work we present a biometric authentication technique based on the combination of multiple biometric readings. The authentication control can be performed offline and the stored identifier does not disclose any information on the biometric traits of the identified person, so that even in case of loss or steal of the document, privacy is guaranteed.

23. Eleni Kosta, Jan Zibuschka, Tobias Scherner, Jos Dumortier: *Privacy-enhancing user-friendly Identity Management for Location Based Services using PRIME technology A legal discussion* [KZSD07].
In *Cyberlaw Security & Private, The Second International Conference on Legal, Security and Privacy Issues in Information Technology (LSPI).*
Abstract. The term Location Based Services (LBS) is used for applications that leverage the users physical location to provide an enhanced service or experience, such as route guidance, tourist and weather information etc. The traditional LBS implementation allows the Mobile Operator as well as the Application Provider of the Location Based Service to have access to a large amount of personal data of the user. The PRIME toolbox, an identity management system, can be used in mobile applications in order to enhance the privacy of the user. In this paper we are going to present the PRIME toolbox and examine it from a legal viewpoint, using a Pollen Warning Application as our case study.

24. Jan Zibuschka, Mike Radmacher, Tobias Scherner, Kai Rannenberg: *Empowering LBS Users: Technical, Legal and Economic Aspects* [ZRSR07].
In *Proc. of the eChallenges conference 2007.*
Abstract. Privacy in computerized environments is perceived very differently depending on the respective point of view. Often privacy enhancing technologies initiated by the user, as a measure of self-defense are seen as conflicting with business goals such as cost-efficiency, revenue assurance, and options for further business development based on existing data. This paper presents the design and implementation of an architecture and prototype for privacy-friendly, interoperable location-based services (LBS), based on intermediation of location data via a location middleware component. The aim is to combine privacy-friendliness, efficiency, and market potential. Therefore the security interests of the stakeholders are analyzed and an architecture solution including an intermediary is introduced. Then the prototype implementation (at a mobile operator) is described and the usage of the prototype for a commercial service and product offer by the operator involved in the development is discussed.

25. Jan Zibuschka, Lothar Fritsch, Mike Radmacher, Tobias Scherner, Kai Rannenberg: *Privacy-Friendly LBS: A Prototype-supported Case Study* [ZFR+07b].
In *Proc. of the 13th Americas Conference on Information Systems.*
Abstract. The development of new products in the mobile data services market poses severe challenges for service providers and mobile network operators. Short-lived products, lack of knowledge about acceptance by users, and the requirement to embed new

products in existing infrastructures lead to difficulties on several levels. Data protection, potentially ambiguous regulation of telecommunication and the existence of a wide range of communication and localization technologies confront product developers with new challenges. Modeling and balancing business interests, data protection requirements, and user preferences while implementing new products on existing infrastructures are some of them. This paper presents the concept and implementation of a prototype that was developed by an international research project with participation of industry partners.

26. Jan Zibuschka, Lothar Fritsch, Mike Radmacher, Tobias Scherner, Kai Rannenberg: *Enabling Privacy of Real-Life LBS: A Platform for Flexible Mobile Service Provisioning* [ZFR+07a].
In *Proc. of the 22nd IFIP TC-11 International Information Security Conference.*
Abstract. Privacy in computerized environments is perceived very differently depending on the respective point of view. Often privacy enhancing technologies initiated by the user as a measure of self-defense, conflict with business goals such as cost-efficiency, revenue assurance, and options for further business development based on existing data. The paper at hand presents the design and implementation of an architecture and prototype for privacy-friendly, interoperable location-based services (LBS), based on intermediation of location data via a location middleware component. The aim is to combine privacy-friendliness, efficiency, and market potential. Therefore the security interests of the stakeholders are analyzed and an architecture solution including an intermediary is introduced. Then the prototype implementation (at a mobile operator) is described and the usage of the prototype for a commercial service and product offer by the commercial operator involved in the development is discussed.

A.3 WP9.0

1. Jan Camenisch, Susan Hohenberger and Michael Østergaard Pedersen. *Batch Verification of Short Signatures.* [CHP07]
In *EUROCRYPT 2007.*
Abstract. With computer networks spreading into a variety of new environments, the need to authenticate and secure communication grows. Many of these new environments have particular requirements on the applicable cryptographic primitives. For instance, several applications require that communication overhead be small and that many messages be processed at the same time. In this paper we consider the suitability of public key signatures in the latter scenario. That is, we consider signatures that are 1) short and 2) where many signatures from (possibly) different signers on (possibly) different messages can be verified quickly. Prior work focused almost exclusively on batching signatures from the same signer.

   We propose the first batch verifier for messages from many (certified) signers without random oracles and with a verification time where the dominant operation is independent of the number of signatures to verify. We further propose a new signature scheme with very short signatures, for which batch verification for many signers is also highly efficient. Combining our new signatures with the best known techniques for batching certificates from the same authority, we get a fast batch verifier for certificates and messages combined. Although our new signature scheme has some restrictions, it is very efficient and still practical for some communication applications.

2. Mira Belenkiy, Melissa Chase, Markulf Kohlweiss and Anna Lysyanskaya. *P-signatures and Noninteractive Anonymous Credentials.* [BCKL08].
   In *Theory of Cryptography Conference (TCC 2008).*
   Abstract. In this paper, we introduce P-signatures. A P-signature scheme consists of a signature scheme, a commitment scheme, and (1) an interactive protocol for obtaining a signature on a committed value; (2) a noninteractive proof system for proving that the contents of a commitment has been signed; (3) a noninteractive proof system for proving that a pair of commitments are commitments to the same value. We give a definition of security for P-signatures and show how they can be realized under appropriate assumptions about groups with a bilinear map. We make extensive use of the powerful suite of non-interactive proof techniques due to Groth and Sahai. Our P-signatures enable, for the first time, the design of a practical non-interactive anonymous credential system whose security does not rely on the random oracle model. In addition, they may serve as a useful building block for other privacy-preserving authentication mechanisms.

3. Jan Camenisch, Gregory Neven and abhi shelat. *Simulatable Adaptive Oblivious Transfer.* [CNas07].
   In *EUROCRYPT 2007.*
   Abstract. We study an adaptive variant of oblivious transfer in which a sender has N messages, of which a receiver can adaptively choose to receive k one-after-the-other, in such a way that (a) the sender learns nothing about the receiver's selections, and (b) the receiver only learns about the k requested messages.

   We propose two practical protocols for this primitive that achieve a stronger security notion than previous schemes with comparable efficiency. In particular, by requiring full simulatability for both sender and receiver security, our notion prohibits a subtle selective-failure attack not addressed by the security notions achieved by previous practical schemes.

4. George Danezis, Claudia Diaz and Carmela Troncoso. *Two-Sided Statistical Disclosure Attack.* [DDT07].
   In *Proceedings of Privacy Enhancing Technologies, 7th International Workshop (PET 2007).*
   Abstract. We introduce a new traffic analysis attack: the Two-sided Statistical Disclosure Attack, that tries to uncover the receivers of messages sent through an anonymizing network supporting anonymous replies. We provide an abstract model of an anonymity system with users that reply to messages. Based on this model, we propose a linear approximation describing the likely receivers of sent messages. Using simulations, we evaluate the new attack given different traffic characteristics and we show that it is superior to previous attacks when replies are routed in the system.

5. Claudia Diaz, Carmela Troncoso and George Danezis. *Does additional information always reduce anonymity?.* [DTD07].
   In *Workshop on Privacy in the Electronic Society 2007.*
   Abstract. We discuss information-theoretic anonymity metrics, that use entropy over the distribution of all possible recipients to quantify anonymity. We identify a common misconception: the entropy of the distribution describing the potentialreceivers does not always decrease given more information.We show the relation of these a-posteriori distributions with the Shannon conditional entropy, which is an average overall possible observations.

6. Len Sassaman and Bart Preneel. *The Byzantine Postman Problem: A Trivial Attack against PIR-based Nym Servers.* [SP07a].
In *Technical Report ESAT-COSIC 2007-001.*
Abstract. Over the last several decades, there have been numerous proposals for systems which can preserve the anonymity of the recipient of some data. Some have involved trusted third-parties or trusted hardware; others have been constructed on top of link-layer anonymity systems or mix-nets.

   In this paper, we evaluate a pseudonymous message system which takes the different approach of using Private Information Retrieval (PIR) as its basis. We expose a flaw in the system as presented: it fails to identify Byzantine servers. We provide suggestions on correcting the flaw, while observing the security and performance trade-offs our suggestions require.

7. Meredith L. Patterson and Len Sassaman. *Subliminal Channels in the Private Information Retrieval Protocols.* [PS07].
In *Proceedings of the 28th Symposium on Information Theory in the Benelux, 2007.*
Abstract. Information-theoretic private information retrieval (PIR) protocols, such as those described by Chor et al. [CGKS95], provide a mechanism by which users can retrieve information from a database distributed across multiple servers in such a way that neither the servers nor an outside observer can determine the contents of the data being retrieved. More recent PIR protocols also provide protection against Byzantine servers, such that a user can detect when one or more servers have attempted to tamper with the data he has requested. In some cases (as in the protocols presented by Beimel and Stahl [BS02]), the user can still recover his data and protect the contents of his query if the number of Byzantine servers is below a certain threshold; this property is referred to as Byzantine-recovery.

   However, tampering with a users data is not the only goal a Byzantine server might have. We present a scenario in which an arbitrarily sized coalition of Byzantine servers transforms the userbase of a PIR network into a signaling framework with varying levels of detectability by means of a subliminal channel [Sim83]. We describe several such subliminal channel techniques, illustrate several use-cases for this subliminal channel, and demonstrate its applicability to a wide variety of PIR protocols.

8. Len Sassaman and Bart Preneel. *Solving the Byzantine Postman Problem.* [SP07b].
In *Technical Report ESAT-COSIC 2007-004.*
Abstract. Over the last several decades, there have been numerous systems proposed which aim to preserve the anonymity of the recipient of some data. Some have involved trusted third-parties or trusted hardware; others have been constructed on top of link-layer anonymity systems or mix networks.

   In this paper, we examine the Pynchon Gate [SCM05], a pseudonymous message system which takes an alternate approach to this problem by using Private Information Retrieval (PIR) as the basis for its pseudonymity properties. We restrict our examination to a flaw in the Pynchon Gate system first described in our technical report [SP07a]; as it was originally presented, the Pynchon Gate detects the presence of (and protects against certain attacks by) Byzantine servers operating in the system, but it fails to identify *which* server or set of servers is Byzantine, thus opening the door for denial of service attacks as well as other potential anonymity compromises by Byzantine servers.

We show a trivial modification to the original PynGP which allows for detection and identification of Byzantine nodes, with no weakening of the security model necessary, at the relatively affordable cost of greater bandwidth requirements during certain communication operations. We demonstrate that this adequately solves the problems raised by [SP07a], and argue that it is the most suitable method of addressing the attack in question yet proposed.

We then evaluate an alternate approach to solving to the problem described in [SP07a], proposed by Goldberg in his recent paper [Gol07]. We compare the security and performance trade-offs made in that proposal, and find it less secure against anonymity attacks as compared to the original (but flawed) Pynchon Gate Protocol (PynGP) [Mat04] presented in the first Pynchon Gate paper. We show that this proposal is significantly weaker than the solution offered in this paper, which retains the security properties of the original Pynchon Gate Protocol.

9. Markulf Kohlweiss, Sebastian Faust, Lothar Fritsch, Bartek Gedrojc and Bart Preneel. *Efficient Oblivious Augmented Maps: Location-Based Services with a Payment Broker.* [KFF$^+$07].
In *Proceedings of Privacy Enhancing Technologies, 7th International Workshop (PET 2007).*
Abstract. Secure processing of location data in location-based services (LBS) can be implemented with cryptographic protocols. We propose a protocol based on oblivious transfer and homomorphic encryption. Its properties are the avoidance of personal information on the services side, and a fair revenue distribution scheme. We discuss this in contrast to other LBS solutions that seek to anonymize information as well as possible towards the services. For this purpose, we introduce a proxy party. The proxy interacts with multiple services and collects money from subscribing users. Later on, the proxy distributes the collected payment to the services based on the number of subscriptions to each service. Neither the proxy nor the services learn the exact relation between users and the services they are subscribed to.

10. George Danezis and Claudia Diaz. *A Survey of Anonymous Communication Channels.* [DD08].
In *Microsoft Technical Report (MSR-TR-2008-35).*
Abstract. We present an overview of the field of anonymous communications, from its establishment in 1981 from David Chaum to today. Key systems are presented categorized according to their underlying principles: semi-trusted relays, mix systems, remailers, onion routing, and systems to provide robust mixing. We include extended discussions of the threat models and usage models that different schemes provide, and the trade-offs between the security properties offered and the communication characteristics different systems support.

11. George Danezis and Len Sassaman.*How to Bypass Two Anonymity Revocation Schemes.* [DS08].
In *Privacy Enhancing Technologies Symposium (PETS 2008).*
Abstract. In recent years, there have been several proposals for anonymous communication systems that provide intentional weaknesses to allow anonymity to be circumvented in special cases. These anonymity revocation schemes attempt to retain the properties of strong anonymity systems while granting a special class of people the ability to selectively

break through their protections. We evaluate the two dominant classes of anonymity revocation systems, and identify fundamental flaws in their architecture, leading to a failure to ensure proper anonymity revocation, as well as introducing additional weaknesses for users not targeted for anonymity revocation.

12. Klaus Kursawe, Peter Palfrader and Len Sassaman. *Echolot and Leuchtfeuer: Measuring the Reliability of Unreliable Mixes.* [KPS07].
In *Technical Report ESAT-COSIC 2007-005.*
Abstract. In a mix-net, information regarding the network health and operational behavior of the individual nodes must be made available to the client applications so they may select reliable nodes to use in each messages path through the mix-net. We introduce the concept of a pinger, an agent which tests the reliability of individual mixes in the mix-net, and publishes results for the mix clients to evaluate. We discuss the security concerns regarding pingers, including the issues regarding anonymity set preservation, information disclosure, and node cheating. We present our software Echolot, the most comprehensive and widely adopted pinger for the Mixmaster anonymous remailer network. To address a serious anonymity weakness potentially introduced by the careless deployment of pingers, we present Leuchtfeuer, a new protocol enhancement for mix-nets. Leuchtfeuer eliminates the active and passive intersection attacks that are possible when different users obtain conflicting reliability statistics about the mix-net.

13. Len Sassaman. *The Faithless Endpoint: How Tor puts certain users at greater risk.* [Sas07].
In *Technical Report ESAT-COSIC 2007-003.*
Abstract. We demonstrate that the decision to employ certain security solutions must be balanced against the additional risks that these security solutions introduce for a given users situation, in the context of its specific threats. As an example case, we consider the anonymity network Tor and examine scenarios where the use of Tor decreases ones overall security. We then show that such trade-offs are reasonable for some, but not all, potential users of Tor. We then consider possible ways to mitigate these risks.

14. Mina Deng and Bart Preneel. *On secure and anonymous buyer-seller watermarking protocol.* [DP08].
In *International Conference on Internet and Web Applications and Services.*
Abstract. Buyer-seller watermarking protocols incorporates digital watermarking with cryptography, in order to protect applications. digital copyrights and privacy rights for the seller and the buyer before, during, and after purchase activities in e-commerce. In this paper, we analyze the security of some previously proposed schemes, and propose an secure and anonymous buyer-seller watermarking protocol. In contrast to early work, our improvement on the protocols security properties ensures that the design requirements are fulfilled. The proposed scheme is able to simultaneously resolve the piracy tracing problem, the customers rights problem, the unbinding problem, the anonymity problem, the conspiracy problem, and the dispute problem. In the proposed scheme, a buyer can purchase digital contents anonymously but his anonymity can be revoked as soon as he is adjudicated to be guilty by a legal institute, such as civil court.

15. George Danezis and Claudia Diaz. *Space-Efficient Private Search.* [DD07].
In *Proceedings of Financial Cryptography (FC2007).*
Abstract. Private keyword search is a technique that allows for searching and retrieving

documents matching certain keywords without reveal- ing the search criteria. We improve the space efficiency of the Ostrovsky et al. Private Search [OS05] scheme, by describing methods that require considerably shorter buffers for returning the results of the search. Our basic decoding scheme recursive extraction, requires buffers of length less than twice the number of returned results and is still simple and highly efficient. Our extended decoding schemes rely on solving systems of simultaneous equations, and in special cases can uncover documents in buffers that are close to 95% full. Finally we note the similarity between our decoding techniques and the ones used to decode rateless codes, and show how such codes can be extracted from encrypted documents.

A.4 WP10.0

1. Leonardo Martucci, Markulf Kohlweiss, Christer Andersson, Andriy Panchenko: *Self-certified Sybil-Free Pseudonyms: Introducing Privacy in Infrastructureless Wireless Networks* [MKAP08].
In *Proceedings of the First ACM Conference on Wireless Network Security (WiSec 2008).*
Abstract. Accurate and trusted identifiers are a centerpiece for any security architecture. Protecting against Sybil attacks in a privacy-friendly manner is a non-trivial problem in wireless infrastructureless networks, such as mobile ad hoc networks. In this paper, we introduce *self-certified Sybil-free pseudonyms* as a means to provide privacy-friendly Sybil-freeness without requiring continuous online availability of a trusted third party. These pseudonyms are self-certified and computed by the users themselves from their cryptographic long-term identities. Contrary to identity certificates, we preserve location privacy and improve protection against some notorious attacks on anonymous communication systems.

2. Christer Andersson, Andriy Panchenko: *Practical Anonymous Communication on the Mobile Internet using Tor* [AP07].
In *Proceedings of the Third International Workshop on the Value of Security through Collaboration (IEEE SECOVAL 2007).*
Abstract. This paper proposes and evaluates several archi- tectural designs for enabling anonymous browsing on the mobile Internet. These architectural designs make use of the Tor network in a mobile setting for the provisioning of anonymity to mobile devices. We compare several architectural designs with respect to their anonymity and performance properties. In particular, we are interested in finding a trade-off between anonymity and performance. We also evaluate the architectural designs against other criteria such as practicality, usability, availability, and trust. We show that the most preferable option given a powerful mobile device and some optimizations in the Tor protocol is the option where the Tor client is run directly on the mobile device.

3. Christer Andersson, Markulf Kohlweiss, Leonardo Martucci, Andriy Panchenko: *A Self-certified and Sybil-Free Framework for Secure Digital Identity Domain Buildup* [AKMP08].
In *Proceedings of the Workshop in Information Security Theory and Practices 2008: Smart Devices, Convergence and Next Generation Networks (WISTP 2008).*
Abstract. An attacker who can control arbitrarily many user identities can break the security properties of most conceivable systems. This is called a “Sybil attack”. We present a solution to this problem that does not require online communication with a trusted third party and that in addition preserves the privacy of honest users. Given an

initial so-called Sybil-free identity domain, our proposal can be used for deriving Sybil-free unlinkable pseudonyms associated with other identity domains. The pseudonyms are self-certified and computed by the users themselves from their cryptographic long-term identities.

4. Lexi Pimenidis, Eleni Kosta: *The impact of the retention of traffic and location data on the internet user* [PK08].
   In *Datenschutz und Datensicherheit, Zeitschrift für rechts- und prüfungssicheres Datenmanagement.*
   Abstract. (this journal article does not have an abstract)

5. Dogan Kesdogan, Vinh Pham, Lexi Pimenidis: *Analyse der Verkettbarkeit in nutzergesteuertem Identitätsmanagement* [KPP08].
   In *Proceedings of Sicherheit 2008.*
   Abstract. Nutzergesteuerte Identitätsmanagementsysteme haben als Ziel die Verkettbarkeit zwischen den verschiedenen digitalen Identitäten eines Benutzers zu verhindern. Wir führen theoretische und experimentelle Untersuchungen zu dem folgenden Informationsverlustproblem durch: Seien konsistente Beobachtungen eines starken Identitätsmanagementsystems gegeben (z.B. *Idemix*), in dem Benutzer Credentials anfordern und vorzeigen. Geben diese Beobachtungen genügend Informationen, um die Verkettungen der digitalen Identitäten der Benutzer aufzudecken?

   Wir werden zeigen, dass die Pseudonyme der Benutzer theoretisch verkettbar sind und dass es ein NP-vollstndiges Problem darstellt. Des Weiteren evaluieren wir praktische Instanzen des Problems und zeigen, dass die digitalen Pseudonyme eindeutig verkettbar sind, trotz vollstndiger Anonymität auf der Netzwerkebene und der Benutzung eines Identitätsmanagementsystems.

6. Andriy Panchenko, Lexi Pimenidis, Hannes Renner: *Performance Analysis of Anonymous Communication Channels Provided by Tor* [PPR08].
   In *Proceedings of the Third International Conference on Availability, Reliability and Security, ARES 2008.*
   Abstract. Providing anonymity for end-users on the Internet is a very challenging and difficult task. There are currently only a few systems that are of practical relevance for the provision of low-latency anonymity. One of the most important to mention is the Tor network that is based on onion routing. Practical usage of the system often leads to delays which are not tolerated by the average end-user. This, in return, discourages many of them from the use of such systems and hence indirectly lowers the protection of remaining users due to a smaller user base. In this paper we show to which extend overloaded nodes and links, as well as geographical diversity of nodes have an influence on the general performance of Tor communication channels. After that, we propose new methods of path selection for performance-improved onion routing which are based on actively measured latencies and estimated available capacities using passive observations of link-wise throughput.

7. Dogan Kesdogan and Vinh Pham and Lexi Pimenidis: *Information Disclosure in Identity Management* [KPP07].
   In *Proceedings of 12th Nordic Workshop on Secure IT-Systems.*
   Abstract. User Controlled Identity Management Systems have the goal to hinder the linkability between the different digital identities of a user. We perform a theoretical and

an experimental study of the following information leakage problem: given a consistent view on the actions of a strong identity management system (e.g. *Idemix*) where k users pseudonymously issue and show some credentials, do these observations disclose sufficient information about the linkability of the digital identities?

We show that in theory, linking the different pseudonyms of a user is a NP-complete problem, by using first order logic. In addition, we evaluate practical instances of the problem, and show that there are non-negligible probabilities that despite full anonymization and use of an identity management system, pseudonyms are unambiguous linkable.

8. Andriy Panchenko and Lexi Pimenidis: *Crowds Revisited: Practically Effective Predecessor Attack* [PP07a].
In *Proceedings of 12th Nordic Workshop on Secure IT-Systems.*
Abstract. Crowds is a peer-to-peer system for protecting users' anonymity for web transactions. One of the more serious disadvantages of it is the degree of anonymity provided with respect to the colluding system members: the one who forwards a message to a colluding node is more likely to be the originator of the message than any other member in the system. Furthermore, with the system size growth, the probability that the request came from the initiator of the communication becomes more likely.

   In this paper we want to assess to which degree Crowds is applicable despite these weaknesses. To this end, we calculate the needed number of observations for colluding members in order to determine with arbitrary precision how often some users communicate with an external service. An additional question that will be addressed is the possibility to hamper this degradation of the provided anonymity level by a method for adaptive behavior of honest members.

9. Lexi Pimenidis: *On Anonymit Loss in Open Environments* [Pim07].
In *Proceedings of 12th Nordic Workshop on Secure IT-Systems.*
Abstract. Some anonymizing networks, like e.g. JAP, relay messages to external services, i.e. allow communication to entities that are not participants of the network. In general this is the case if networks provide access to the world wide web (like e.g. JAP, and Tor) or allow email messaging with pseudonymous return addresses (like e.g. Mixmaster).

   Even if answers are indistinguishable from requests in the anonymizing network, they can severly damage the amount of anonymity provided by a system. This is due to the fact that outside the anonymous network answers are linkable to their request, take an individual amount of time to be generated, and thus are returned to the original sender in different anonymity sets.

   We will show in this paper that anonymizing networks which provide access to external services are subject to a traffic analysis that can strongly degrade the anonymity provided to their participants. The degradation is based on the fact that replies to messages are linkable outside the anonymizing network. This analysis can be used as an effective measure to preproccess all kinds of anonymized communication and build a step stone for further, more elaborated, attacks.

10. Olaf Landsiedel, Lexi Pimenidis, Klaus Wehrle, Heiko Niedermayer, Georg Carle: *More: A Peer-To-Peer Based Connectionless Onion Router* [LPW$^+$07].
In *Proceedings of IEEE GLOBECOM, Globalcommunications Conference.*
Abstract. Onion Routing is today's typical substrate for anonymous near-real-time communication. Via layered encryption, Onion Routers such as TOR or Tarzan build a

static tunnel through a peer-to-peer relay network. All traffic exchanged between two end points uses one and the same tunnel, making the design susceptible to attacks based on pattern analysis. Recent publications even extend this theoretical threat by showing the practical feasibility of a pattern analysis attack on the deployed TOR system.

In contrast to today's anonymous communication systems, Core routes each packet a different communication path and so is not susceptible to this class of attacks. Furthermore, inspired by IP-routing the connectionless approach reduces the complexity of the Onion Router.

In this paper, we describe the design of our Connectionless Onion Router, evaluate its performance, and address the communication overhead. We present address virtualization to abstract from Internet addresses and to provide transparent application support. Thus, no application-level gateways or proxies are required to sanitize protocols from network level information. Acting as an IP-datagram service, our scheme provides a substrate for anonymous communication to a wide range of applications using TCP and UDP.

11. David Piegdon and Lexi Pimenidis: *Targeting Physically Addressable Memory* [PP07c].
    In *Proceedings of Fourth GI International Conference on Detection of Intrusions & Malware, and Vulnerability Assessment.*
    Abstract. This paper introduces new advances in gaining unauthorised access to a computer by accessing its physical memory via various means. We will show a unified approach for using IEEE1394, also known as firewire, file descriptors and other methods to read from and write into a victim's memory. Thereafter we will show the power of this ability in several example attacks: stealing private SSH keys, and injecting arbitrary code in order to obtain interactive access with administrator privileges on the victim's computer.

    These advances are based on data structures that are required by the CPU to provide virtual address spaces for each process running on the system. These data structures are searched and parsed in order to reassemble pages scattered in physical memory, thus being able to read and write in each processes virtual address space.

    The attacks introduced in this paper are adaptable to all kinds of operating system and hardware combinations.

12. Andriy Panchenko and Lexi Pimenidis: *Using Trust to Resist Censorship in the Presence of Collusion* [PP07b].
    In *Proceedings of 22nd IFIP TC-11 International Information Security Conference.*
    Abstract. Censorship resistance deals with an attempt to prevent censors from the acquaintance of distribution of a particular content through the network. Providing resistance against censoring is a very challenging and difficult task to achieve. However it is vital for the purpose of freedom of speech, mind and achievement of democratic principles in todays society.

    In this paper we define a model of a censorship resistant system. Thereafter we propose to split the problem of resisting censorship into the following two sub-problems: a trusted directory and steganographic data transfer. The directory is used in order to prolong contacts among peers based on their reputation in a way, that honest members get contacts only to other honest peers and colluded members remain isolated. Furthermore, we aim to provide an analysis of a trusted directory for reputation and its implications on censorship resistant systems. To this end we define a set of properties that such a directory has to fulfill and develop a proposal for the implementation. Finally we provide a simulation-based validation of our approach.

A.5 WP11.0

1. Stefan Berthold and Sebastian Clauß: *Linkability Estimation Between Subjects and Message Contents Using Formal Concepts* [BC07].
   In *Proceedings of the 2007 ACM Workshop on Digital Identity Management.*
   Abstract. In this paper, we examine how conclusions about linkability threats can be drawn by analyzing message contents and subject knowledge in arbitrary communication systems. At first, we define messages described by their contents as formal contexts. Then, we define subjects described by their knowledge as further formal contexts. Finally, we show that concept lattices, which are achieved by applying Formal Concept Analysis to the concatenation of these formal contexts, can be used in order to draw conclusions about correlations, and therefore linkability, between contents of messages and knowledge of subjects. The goal is to define formal specifications which can be utilized in privacy enhancing identity management systems in order to support users in the choice of data items which are to be disclosed to a communication partner.

2. Sebastian Clauß: *Towards Quantification of Privacy Within a Privacy-Enhancing Identity Management System* [Cla07].
   PhD Thesis, Technische Universität Dresden, 2007.
   Abstract. As the name implies, privacy-enhancing identity management shall empower users to use services provided by the Internet in a privacy-friendly way. In order to achieve this goal, the user is supported in minimising disclosure of personal identifiable information (PII), while preserving accountability when necessary. Major techniques for that are using digital pseudonyms and various cryptographic mechanisms which enable to achieve confidentialityas well as integrity- and accountability-related protection goals.

   Because in many cases disclosure of at least some PII is necessary, it is desirable for a user to at least have an overview about data known by certain communication partners and which is the main goal of this work to measure to what extent these data can be used to identify a person, i.e. to measure the users anonymity.

   Another major goal in order to retain users privacy is to prevent or at least make it difficult to create user profiles containing a large amount of data. Even if it may not be directly possible to identify the person behind a profile, the larger the amount of data in the profile, the more the privacy of the corresponding user will be sacrificed in case additional data enables identification. So, profiles should be kept small, i.e., if not necessary for providing a service, different actions should be decoupled in a way that it is impossible or at least difficult to assign PII disclosed in different actions to the same profile. Therefore, a users actions should be kept unlinkable.

   This work deals with opportunities and drawbacks of measuring anonymity and linkability of actions in the scope of a privacy-enhancing identity management system (PE-IMS). Thereto, we first give a detailed characterisation of a PE-IMS goals and mechanisms.

   We show how anonymity metrics already known from evaluating anonymity providing services on the network layer can be used at application layer. We develop a model which can be used to measure anonymity and linkability of actions based on statistical knowledge gained from disclosure of user data during actions.

We evaluate efficiency of measurement calculations and propose situationdependent improvements. With regard to practical applicability, we evaluate the influence of misinformation and incomplete information on the measurement results.

We discuss how results of anonymity and linkability measurements can practically be used in the scope of PE-IMS. Therefore we propose possibilities for direct information of the user as well as possibilities for (semi-automatic) usage within disclosure policies.

Finally, we evaluate the procedure of measuring anonymity and linkability with respect to the multilateral security model of PE-IMS. It turns out that within this security model, anonymity measurement performed by single users themselves is risky, because for a comprehensive measurement each user would need data invading the privacy of others to a large extent. Therefore we discuss possibilities to have third parties performing measurements, and we evaluate (cryptographic) mechanisms to secure the privacy-relevant data as well as the measurement calculations in an appropriate manner.

3. Andreas Pfitzmann, Andreas Juschka, Anne-Katrin Stange, Sandra Steinbrecher, Stefan Köpsell: *Communication Privacy* [PJS+08].
In *Digital Privacy: Theory, Technologies, and Practices, Auerbach Publications.*
Abstract. Many people have a fallacious feeling of being anonymous when surfing the Internet. But ordinary Internet communication on the network layer is by default not anonymous because of the usage of identifying characteristics like IP or MAC address. So if no additional measures are taken, an adversary can easily observe which participants of a network communicate with each other. But anonymity on the network layer of communication systems can be achieved by the use of anonymizing techniques. This book chapter presents an overview of anonymizing techniques which enable anonymity in a communication network.

4. André Adelsbach, Ulrich Greveler, Stephan Groß, Sandra Steinbrecher: *ANOCAST: Rethinking Broadcast Anonymity in the Case of Wireless Communications* [AGGS08].
In *Proceedings of Sicherheit 2008.*
Abstract. In this work we present the ANOCAST environment being an anonymous communication framework based on wireless technology having reached prototype status. By exploiting the availability of current broadcast communication systems (e.g. digital satellites and Wi-Fi) we achieve recipient anonymity and unobservability. We examine our approach by a simulation approach using a prototype implementation, raising several questions concerning its scalability, efficiency and possible application scenarios as well as a discussion on its security accomplishments.

   The key finding of our work shows that practicability of anonymous communication for the masses can be realised today and the community should rethink the common consensus that enormous traffic overhead jeopardises any practical anonymous communication system.

5. Katrin Borcea-Pfitzmann, Marit Hansen, Katja Liesebach, Andreas Pfitzmann and Sandra Steinbrecher: *Managing Ones Identities in Organisational and Social Settings* [BPHL+07].
In *Datenschutz und Datensicherheit.*
Abstract. While interacting over the Internet, users should be empowered to reveal only those subsets of their personal attributes - called partial identities - which are appropriate for the actual situation and context. Avoiding the use of few and easily linkable partial identities is a prerequisite for privacy protection. Traditional user-controlled identity

management systems primarily support individuals interacting with organisations, but generally ignore special needs which arise if individuals interact with each other. In order to support online communities, those systems have to change.

6. Simone Fischer-Hübner, John Sören Pettersson, Mike Bergmann, Marit Hansen, Siani Pearson, Marco Casassa Mont: *HCI Designs for Privacy-Enhancing Identity Management* [FHPB+08].
Chapter 11 in *Digital Privacy: Theory, Technologies, and Practices.*
(Joint work with WP6.1 and WP12.0, please refer to Section A.1 for the abstract.)

7. Marit Hansen, Simone Fischer-Hübner, John Sören Pettersson, Mike Bergmann: *Transparency Tools for User-controlled Identity Management* [HFHPB07].
In *Expanding the Knowledge Economy: Issues, Applications, Case Studies; Proceedings of eChallenges 2007.*
(Joint work with WP6.1, please refer to Section A.1 for the abstract.)

8. Marit Hansen, *Marrying Transparency Tools with User-Controlled Identity Management* [Han08a].
In: *"The Future of Identity in the Information Society": Proceedings of the Third IFIP WG 9.2, 9.6/11.6, 11.7/FIDIS International Summer School.*
(Joint work with WP6.1)
Abstract. User-controlled identity management systems assist individuals in managing their private sphere. An individuals privacy can be supported by transparency on processing of personal data. After giving an overview on transparency properties as well as its relation to privacy and data protection regulation, this text introduces different transparency tools: Prior to an interaction, information on the interacting party should be made transparent. During the interaction, privacy policies have to be communicated. Afterwards, users should be helped in exercising their privacy rights such as, among others, the right to access own personal data. In addition information on security and privacy incidents provides complementary data for the users perception of the level of privacy. Although transparency tools alone are no panacea for maintaining the private sphere, the combination of transparency tools and user-controlled identity management systems yields viable functionality to empower users to protect their privacy.

9. Marit Hansen: *User-controlled identity management: key to the future of privacy?* [Han08b].
In *International Journal of Intellectual Property Management (IJIPM), Special Issue on: "Identity, Privacy and New Technologies".*
Abstract. Processing personal data is a prerequisite of todays participation in the Information Society which increasingly implies threats to individuals privacy. Classical identity management systems may even intensify those privacy threats. However, comprehensive user-controlled identity management systems take a different approach as they implement privacy principles and aim at helping users in managing their privacy. This text distinguishes various notions of privacy and identifies linkability control as a key concept underlying all privacy issues. The described building blocks for user-controlled identity management support privacy in different ways, e.g., by separating contexts, workflows and the related partial identities from each other, by stating and enforcing privacy policies governing also linkability, or by enhancing transparency to inform the user about potential linkages. User-controlled identity management systems are not a panacea for

all privacy problems, but they can act as a gateway and guardian for individuals privacy. By this, they have the potential to become a key tool for future privacy concepts.

10. Ellen Tweney, Stephen Crane: *Trustguide2: An exploration of privacy preferences in an online world* [TC07].
In: *eChallenges 2007.*
Abstract. Today's world is dominated by on-line services that request personal information from individuals. Concerns arise when it is not clear to the individual why such information is required or how it will be used. This lack of control leads to suspicion and ultimately distrust. Returning control to the individual by permitting them to state preferences for the management of their personal information is one approach to reversing this trend. From previous studies we know that control engenders trust. In Trustguide2 we explore, through a series of user engagements, the form that these controls might take in relation to managing personal identifying information.

11. Elisabetta Carrara and Giles Hogben (Joint Ed.): *Reputation-based Systems: a security analysis* [CH07].
ENISA Position Paper No. 2, (contributed to this paper).
Abstract. Reputation allows users to form an expectation of behaviour based on the judgements of others, bringing the significant economic and social benefit of being able to trust others who are not known directly. Reputation can encourage good behaviour, as users seek good reputation and benefit from it. It follows that electronic reputation is becoming as valuable an asset as traditional reputation. As new applications embrace reputation-based systems, the value of online reputation will continue to grow — and will become increasingly the target of attacks. This paper explains the main characteristics of electronic reputation systems and the security related benefits they can bring when used within applications and services. Four main use-cases are described: online markets, peer-to-peer networks, anti-spam techniques and public key authentication (web-of-trust). From these, the main threats and attacks against reputation systems are derived, as well as the security requirements for system design. This leads to a set of core recommendations for best practices in the use of reputation systems.

A.6 WP12.0

1. Marco Casassa Mont, Filipe Beato: *On Parametric Obligation Policies: Enabling Privacy-aware Information Lifecycle Management in Enterprises* [MB07d].
In: *HPL Technical Report, HPL 2007-7.*
Abstract. Enterprises that collect and process personal data must deal with related privacy management issues. It is not just a matter of privacy-aware access control: privacy obligation policies, dictating duties and expectations on how personal data has to be handled, must be considered too. The management of obligation policies is a promising area (affecting the lifecycle management of personal data) but it is still underestimated. Enterprises require solutions that enable automation and leverage their current identity management solutions. HP Labs have been working on this topic in the last few years, also in the context of the EU PRIME project. In this paper we present our recent work on parametric obligation policies and a related obligation management framework to deal with a scalable management of these obligation policies on large amounts of data, stored in distributed data repositories.

2. Marco Casassa Mont, Filipe Beato: *On Parametric Obligation Policies: Enabling Privacy-aware Information Lifecycle Management in Enterprises* [MB07c].
In: *Proceedings of the 8th IEEE International Workshop on Policies for Distributed Systems and Networks 2007.*
Abstract. Enterprises that collect and process personal data must deal with related privacy management issues. It is not just a matter of privacy-aware access control: privacy obligation policies, dictating duties and expectations on how personal data has to be handled, must be considered too. The management of obligation policies is a promising area but it is still underestimated. Enterprises require solutions that enable automation and can leverage their current identity management solutions. HP Labs have been working on this topic in the last few years, also in the context of the EU PRIME project. In this paper we present our recent work on parametric obligation policies and a related obligation management framework to deal with a scalable management of these policies on large amounts of data, stored in distributed data repositories.

3. Adrian Baldwin, Marco Casassa Mont, Yolanta Beres, Simon Shiu: *On Identity Assurance in the Presence of Federated Identity Management Systems* [BMBS07].
In: *Proceedings of the DIM'07 – 207 ACM Workshop on Digital Identity Management.*
Abstract. In this paper we address the appropriate management of risk in federated identity management systems by presenting an identity assurance framework and supporting technologies. We start by discussing the risk mitigation framework that should be part of any identity assurance solution. We then demonstrate how our model based assurance technologies can be used to report success of an identity assurance programme. We discuss how this approach can be used to gain trust within a federated identity management solution both by communicating the nature of the assurance framework and that risks are successfully being mitigated. Finally, we show the importance of automation of controls in easing operational costs (and we describe related approaches developed at HP Labs and PRIME project); providing improved audit information and changing the risk mitigation landscape.

4. Adrian Baldwin, Marco Casassa Mont, Simon Shiu: *On Identity Assurance in the Presence of Federated Identity Management Systems* [BMS07].
In: *Proceedings of the DIM'07 - 207 ACM Workshop on Digital Identity Management.*
Abstract. In this paper we address the appropriate management of risk in federated identity management systems by presenting an identity assurance framework and supporting technologies. We start by discussing the risk mitigation framework that should be part of any identity assurance solution. We then demonstrate how our model based assurance technologies can be used to report success of an identity assurance programme. We discuss how this approach can be used to gain trust within a federated identity management solution both by communicating the nature of the assurance framework and that risks are successfully being mitigated. Finally, we show the importance of automation of controls in easing operational costs; providing improved audit information and changing the risk mitigation landscape.

5. Marco Casassa Mont, Boris Balacheff: *On Device-based Identity Management in Enterprises* [MB07a].
In: *HPL Technical Report, HPL 2007-53.*
Abstract. This paper focuses on the management of device-based identities within enterprises. This is a key requirement in enterprises where the identities of devices have become

as important as the identities of humans (users) to grant access to enterprise resources. In this context, access control systems need to understand which devices are being used to access resources, by whom and in which contexts. Trust in managed devices' identities is an important first step to enable this. Most related commercial solutions are deployed at the network level. Instead, we focus at the application/service level to leverage current enterprise identity management solutions, used to manage users' identities. We investigate requirements and related issues. We introduce an initial approach and describe our related solution. A working prototype (proof-of concept) has been fully implemented by extending HP OpenView Identity Management solutions and using trusted computing-enabled devices. This is work in progress: we aim at setting the context and discussing our current status and next steps.

6. Marco Casassa Mont, Boris Balacheff: *On Device-based Identity Management in Enterprises* [MB07b].
   In: *Trust, Privacy and Security in Digital Businesses.*
   Abstract. This paper focuses on the management of device-based identities within enterprises. This is a key requirement in enterprises where the identities of platforms and devices have become as important as the identities of humans to grant access to enterprise resources. In this context, access control systems need to understand which devices with what properties are being used to access resource, by whom and in which contexts. Trust in managed devices' identities is an important first step to enable this. No effective commercial solution is currently available. We investigate requirements and related issues. We introduce an initial approach to: model devices' identities; enable their provisioning in heterogeneous enterprise systems; provide support for making and enforcing related access control decisions; leverage trusted computing capabilities of modern devices to deal with aspects of trust management. We describe a related solution where access control is based on policies that take into account: device identities in addition to traditional human-based identities; protected resources; additional constraints on contextual information. A working prototype (proof-of concept) has been fully implemented by HP Labs by leveraging and extending HP OpenView Identity Management solutions and using trusted computing- enabled devices. This is work in progress: we aim at setting the context and discussing our current status and next steps.

7. Siani Pearson, Marco Casassa Mont, Manny Novoa: *Securing Information Transfer within Distributed Computing Environments* [PMN07].
   In: *HPL Technical Report, HPL 2007-10.*
   Abstract. Today, personal and confidential information is stored on a broad variety of enterprise data resources. This data needs to be secured and protected. Technology must adapt to mitigate the new threats and risks arising from the trend towards dynamic enterprises, adaptive data centres and on-demand allocation of resources. Applications and services are becoming mobile across multiple resources, sometimes in a dynamically allocated way, necessitating migration of sensitive and private data. In this article we propose a policy-driven data protection system to address the inadequacies of current technological solutions in preserving the confidentiality and privacy of data whilst it is migrated between platforms. More specifically, we describe our solution for securing migration of credentials that we are developing for productization.

8. Siani Pearson, Marco Casassa Mont, Manny Novoa: *Securing Information Transfer within Distributed Computing Environments* [PMN08].

In: *IEEE Security & Privacy Magazine.*
Abstract. Today, personal and confidential information is stored on a broad variety of enterprise data resources. This data needs to be secured and protected. Technology must adapt to mitigate the new threats and risks arising from the trend towards dynamic enterprises, adaptive data centres and on-demand allocation of resources. Applications and services are becoming mobile across multiple resources, sometimes in a dynamically allocated way, necessitating migration of sensitive and private data. In this article we propose a policy-driven data protection system to address the inadequacies of current technological solutions in preserving the confidentiality and privacy of data whilst it is migrated between platforms. More specifically, we describe our solution for securing migration of credentials that we are developing for productization.

9. Marco Casassa Mont: *Automation of Privacy Management* [Mon07].
In: *Digma Magazine.*
Abstract. Enterprises need to collect personal information, digital identities and user profiles (from customers, employees and third party partners) to enable their business processes and to customize the provision of their services. This is not going to change in the foreseeable future. However, various privacy laws, in different countries, including EU Data Protection Acts, HIPPA, COPPA, Gramm-Leach-Bliley Act, Japanese Privacy Laws, etc., dictate criteria, principles and constraints on how personal data should be collected, processed and disclosed. Thus, enterprises are coming under increasing pressure to improve their privacy management practice, to satisfy customers and to comply with both regulation and internal guidelines. However, their current approach to privacy management is mainly based on human processes which are expensive and prone to failure - the scale of this problem highlights the desire for additional technology to be part of the solution. The trend towards complexity and dynamism in system configurations further heightens the need for automation to ensure that privacy and security properties are maintained as changes occur, and in addition to check that the delivery of privacy is operating correctly. R&D work has been done at HP Labs to address identity and privacy management problems in this space.

10. Siani Pearson, Marco Cassasa Mont and Stephen Crane: *Analysis of Trust Properties and Related Impact of Trusted Platforms* [PMC07].
In: *Trust Management.*
Abstract. This paper draws a distinction between persistent and dynamic trust and analyses this distinction within the context of trusted computing technology.

11. Stephen Crane, Marco Cassasa Mont and Siani Pearson: *On Helping Individuals to Manage Privacy and Trust* [CMP07].
In: *Trust Management.*
Abstract. Being able to say with absolute certainty that another party can be trusted to handle personal information with today's technology is probably unrealistic. In this paper we explain an approach to establishing trust based on the status of a remote platform and an anticipated willingness of the other party to comply with prior negotiated obligations. Ongoing monitoring and notification, and the ability of the individual to form a simple record of past interaction, provides the individual with greater confidence in situations where they need to share personal sensitive information with organizations they would otherwise not be able to claim they trust. We describe the principles of our approach and architectures that support a practical implementation.

12. Tariq Ehsan Elahi and Siani Pearson: *Privacy Assurance: Bridging the Gap Between Preference and Practice* [EP07a].
In: *Trust, Privacy and Security in Digital Business.*
Abstract. Personal identifying information is released without much control from the end user to service providers. We describe a system to scrutinize the stated claims of a service provider on safeguarding PII by interrogating their infrastructure. We attempt to empower end users by providing means to communicate their privacy concerns in a common language understood by the service provider, allowing them to set baseline privacy practices for service providers to adhere to, and providing a means of retrieving information from the service provider in the common language to base their PII release decisions.

13. Simone Fischer-Hübner, John Sören Pettersson, Mike Bergmann, Marit Hansen, Siani Pearson, Marco Casassa Mont: *HCI Designs for Privacy-Enhancing Identity Management* [FHPB^{+}08].
Chapter 11 in *Digital Privacy: Theory, Technologies, and Practices.*
(Joint work with WP6.1 and WP11.0, please refer to Section A.1 for the abstract.)