First Report on Research on Protocols and Tools for Future eID Solutions

D24.1

Document Identification	
Date	November 6, 2013
Status	Final
Version	1.0

Related SP/WP	SP2/WP24	**Document Reference**	D24.1
Related Deliverable(s)	D12.3, D12.4, D22.1, D22.2, D22.3, D23.1, D34.1, D34.2	**Dissemination Level**	PU
Lead Participant	IBM	**Lead Author**	Jan Camenisch (IBM) Alfredo Rial (IBM)
Contributors	Jan Camenisch (IBM) Alfredo Rial (IBM) Thomas Groß (UNEW) Kovila Coopamootoo (UNEW) M. Nuria Ituarte Aranda (ATOS) Daniel Slamanig (TUG) Christian Hanser (TUG) Sebastian Mödersheim (DTU) Omar Almousa (DTU)	**Reviewers**	Detlef Houdeau (INFINEON) Juan Perez Baun (ATOS)

Abstract: This deliverable describes the research conducted by work package 23 during the first year of the FutureID project with its five tasks: extending languages and tools for compositional reasoning (Task 24.1), establishing methods and languages for privacy goals (Task 24.2), development of privacy-friendly audit and data-handling mechanisms (Task 24.3), development of privacy-friendly revocation mechanisms (Task 24.4) and development of methods for usable privacy (Task 24.5).

This document is issued within the frame and for the purpose of the *FutureID* project. This project has received funding from the European Unions Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 318424.

1 Executive Summary

The aim of work package 24 is to address various short comings of existing and emerging eID solutions. In particular, we aim at extending the toolbox for the formal analysis to cope with the challenges arising in this context, such as the modelling and verification of complex, composed protocols and their privacy features. Another focus is on the development of new cryptographic mechanisms and protocols that complement privacy-enhanced credentials to match requirements that arise in large-scale environments. To this end, work package 24 conducts research on the following five tasks, and this deliverable describes the research conducted on these tasks during the first year of the FutureID project and provides a summary of the research results.

Task 24.1: Extending languages and tools for compositional reasoning. In general reasoning about complex identity systems is impossible, both for the human mind and for the formal verification tools. We thus need compositional reasoning, i.e., theorems that allow to infer the security of a system from the security of the components given several sufficient conditions. We also need to be able to express the required assumptions and guarantees of the components in specification languages like our APS language (defined in Deliverable 42.3) and checker tools for said sufficient conditions. It turns out at the core of many compositionality questions is the disjointess of message formats and we have first results in a draft paper status.

Task 24.2: Establishing methods and languages for privacy goals. Most formal verification approaches are trace based, i.e., one tries to find an attack trace of a distributed system or prove that no possible trace is an attack. Previous research in formal verification has concentrated however on the notion of static equivalence of frames, i.e., whether the attacker can distinguish two given worlds. As a first result we have introduced a new notion of privacy goals that links the classical trace based with the static equivalence based approaches. This allows to specify privacy goals in a very declarative way as a formula describing the information deliberately revealed which fits well especially with privacy friendly identity systems, and at the same time can employ existing research results and verification methods.

Task 24.3: Privacy-friendly audit and data-handling mechanisms. Task 24.3 conducts research on data-handling mechanisms and on audits. Data-handling mechanisms determine how user data is managed by the service provider. To facilitate the design of data-handling policies and mechanisms, privacy-preserving authentication and access control should be employed in order to minimize the data that data-handling policies and mechanisms must deal with. Therefore, we have designed several protocols for privacy-preserving authentication and access control: an oblivious transfer with access control protocol that is UC secure [7], a private set intersection protocol that provides fairness [53], UC-secure building-block protocols for circuit evaluation [30], and an optimally private access control protocol in which the service provider only learns whether the user is granted access or not [76]. Moreover, we provide a language to facilitate the implementation and adoption of privacy-preserving authentication protocols [29], and we investigated two use-cases: the Austrian eID ecosystem [95], and the architecture proposed in the Secure Identity Across

Borders Linked (STORK) [96] project.

Additionally, audits are required to ensure that service providers fulfill the pertinent data-handling policies. To enable privacy-preserving audits, we propose mechanisms based on blank digital signatures (BDS) [67] and warrant-hiding proxy signatures (WHPS) [68], and another mechanism based on commitment schemes with efficient selective disclosure [33]. We also show how to extend the Identity Mixer anonymous credential system to enable signatures on committed graphs and explain its application to audits.

Task 24.4: Development of privacy-friendly revocation mechanisms. An important aspect in the trust reputation of identity providers is their ability to react on changes, in particular to revoke credentials in case they get lost or corrupted, or a user lost his right to possess a certain credential. Such revocation is more challenging when advanced identity schemes that support pseudonymous authentication are used, since therein different transactions of the same user are supposed to be unlinkable. There exist already a variety of cryptographic approaches that solves revocation in a privacy-friendly manner [26], [34], [31], [32], [83] but some important challenges remain, such as the (partial) revocation of pseudonyms, or solutions that are non-linear in the number of (revoked) users. We describe ongoing work but still do not have results that achieve those goals.

Task 24.5: Development of methods for usable privacy. Privacy is a key requirement for identity management. However usability of online privacy, that is ensuring individuals are able to control whether, when and who has access to what personal information is a complex task. This is mainly because privacy spans across social, psychological, legal and technical dimensions. The aim of Task 24.5 is to build on the findings of previous research in usable privacy of identity management and to develop new methods of engaging with users. This involves research into mental models of privacy which will help towards better depicting users' privacy in relation to their identity in different context and further contribute to making informed consent decisions. First, we review the findings and recommendations on the usability of privacy enhancing identity management systems from the PRIME and PrimeLife projects. Then we pursue a case study with privacy in real-world environments and finally we lay the foundations for new insights in usability for privacy by helping users establish appropriate mental models.

2 Document information

2.1 Contributors

Name	**Partner**
Jan Camenisch	IBM
Alfredo Rial	IBM
Thomas Groß	UNEW
Kovila Coopamootoo	UNEW
M. Nuria Ituarte Aranda	ATOS
Daniel Slamanig	TUG
Christian Hanser	TUG
Sebastian Mödersheim	DTU
Omar Almousa	DTU

2.2 History

0.01	2013-09-24	Alfredo Rial	1st Draft
0.02	2013-10-17	Thomas Groß	1st Draft Compositional Reasoning Task 24.1
0.03	2013-10-22	Alfredo Rial	1st Draft Executive Summary
0.04	2013-10-22	Alfredo Rial	1st Draft Introduction
0.05	2013-10-22	Alfredo Rial	Abstracts of IBM papers
0.06	2013-10-22	Kovila Coopamootoo and Alfredo Rial	Task 24.5
0.07	2013-10-22	Sebastian Mödersheim	Task 24.2
0.08	2013-10-23	Daniel Slamanig	Task 24.3 TUG papers
0.09	2013-10-23	Alfredo Rial	Task 24.3 IBM papers
0.10	2013-10-23	Kovila Coopamootoo and Alfredo Rial and Sebastian Mödersheim	Executive Summary
0.11	2013-10-23	Alfredo Rial	Conclusion
0.12	2013-10-23	Alfredo Rial	Task 24.4
0.13	2013-10-23	Thomas Groß	Task 24.1
0.14	2013-10-23	Daniel Slamanig	progress of Task 24.4

2.3 Table of Contents

2.4 List of Acronyms

ABC	Attribute-Based Credential
ACM	Association for Computer Machinery
ACOT	Oblivious Transfer with Access Control
AIF	AVISPA Intermediate Format
APS	Authentication Protocol Specification
AVISPA	Automated Validation of Internet Security Protocols and Applications
BDS	Blank Digital Signatures
CARL	Card-Based Access Control Requirements Language
CRS	Common Reference String
DDH	Decisional Diffie-Hellman
DSA	Digital Signature Algorithm
eID	Electronic Identification
G3C	Graph 3-Colorability
HCI	Human-Computer Interaction
IP	Internet Protocol
LPAR	Logic for Programming Artificial Intelligence and Reasoning
MOA	Modules for Online Applications
NP	Non-deterministic Polynomial Time
OFMC	On-the-fly Model Checker
OT	Oblivious Transfer
PACE	Password-Authenticated Channel Establishment
PEPS	Pan-European Proxy Service
PET	Privacy-Enhancing Technology
PKIX	Public-Key Infrastructure X.509
ProVerif	Cryptographic protocol verifier
PSI	Private Set Intersection
RSA	Rivest Shamir Adleman
SAML	Security Assertion Markup Language
SEMIRAMIS	Secure Management of Information across multiple Stakeholders
SPASS	SPASS: An Automated Theorem Prover for First-Order Logic with Equality
STORK	Secure identity across borders linked
SWIFT	Secure Widespread Identities for Federated Telecommunications
SXDH	Symmetric external Diffie-Hellman
TLS	Transfer Layer Security
UC	Universal Composability
USES	Users' Self-Estimation Scale
WHPS	Warrant-Hiding Proxy Signatures
WP	Work Package
XDLIN	External Decision Linear
XML	Extensible Markup Language

2.5 Glossary of Terms

access control
Prevention and protection of resources against unauthorised access; a process by which use of resources is regulated according to a security policy and is permitted by only authorised people according to that policy.

Access control can be logical e.g. for IT systems, or physical e.g. for entry to buildings.

composition
Combining protocols, components or sub-systems to larger systems.

compositionality
Property of components and systems that guarantees that they can be composed securely, maintaining their properties. Roughly equivalent to composability, usually used in the formal methods context.

composability
Property of components and systems that guarantees that they can be composed securely, maintaining their properties. Roughly equivalent to compositionality, usually used in the cryptography context. Notable variants are Reactive Simulatability and Universal Composability, stating that a component can be composed with an arbitrary environment and will still maintain its properties.

parallel composition
Protocol composition in which several protocols are executed over the same communication medium. The protocols are possibly using the same key-infrastructure.

sequential composition
Protocol composition in which one protocol is executed after another, the output of the first one feeding as input into the second one.

vertical composition
Protocol composition in which one protocol is executed on top of another protocol, for instance, as an application protocol is run over a secure channel.

3 Introduction

The objective of work package 24 is to develop protocols and tools that can be used in the next generation of privacy-enhanced eID solutions. In particular, we aim at extending the toolbox of formal methods in order to cope with the challenges arising in this context, such as the modeling and verification of complex, composed protocols and their privacy features. Another focus is on the development of new cryptographic mechanisms and protocols that complement privacy-enhanced credentials and improve their supported functionalities in large-scale environments. To this end, work package 24 conducts research on the following tasks:

Task 24.1: Extending languages and tools for compositional reasoning

Task 24.2: Establishing methods and languages for privacy goals

Task 24.3: Development of privacy-friendly audit and data-handling mechanisms

Task 24.4: Development of privacy-friendly revocation mechanisms

Task 24.5: Development of methods for usable privacy

In the following subsections, we describe the research goals of each of the tasks.

3.1 Extending Languages and Tools for Compositional Reasoning (Task 24.1)

Task 24.1 aims at providing languages and tools that allow the analysis of complex systems that are composed of multiple components. This is useful for the eID based solutions developed in FutureID that we want to formally analyze with the OFMC [19], AIF/ProVerif/SPASS toolchain [82], [22], [94], which are complex systems. They are composed from several smaller components, such as secure channel protocols and application protocols that are run over such a channel, as well as compositions of eID protocols with identity federation and anonymous credential systems. The direct verification of composed systems is often too complex and also not desirable, because any change of the composition will invalidate the overall security statement. In order to verify systems compositionally, we will establish suitable interfaces and abstract properties, extend existing work on protocol composition [47], [65], [43], [42], [64] with channels suitable for eID protocols (e.g., with unauthenticated or unilateral authenticated end-points) as well as the development of extended specification languages and tools.

The goals of this work are two-fold: First, we develop a set of design principles, such as disjointness of the message format when messages have a different meaning. These good engineering practices [6] avoid many problems by construction already. Second, we aim at establishing compositionality theorems that show that systems that adhere to the identified design principles and that are safe in isolation can be arbitrarily composed without introducing new vulnerabilities.

3.2 Establishing Methods and Languages for Privacy Goals (Task 24.2)

Task 24.2 aims at establishing methods and languages for the analysis of privacy goals with formal methods tools. It pursues that goal with two sub-tasks, one to establish a privacy analysis method and its formalization, and the other to establish semantics for claims languages to allow reasoning over them.

Task 24.2.1: Formal Methods for Privacy Goals. The protection of personal data is very important but often neglected in formal analysis for its intricacy: it is not sufficient to evaluate single runs of a system (as for classical secrecy properties for instance) but one must at least consider whether an intruder can distinguish several runs of a system. This provides the basis for formulating many privacy properties; further results can be achieved by consideration of context and his auxiliary knowledge. Currently there is practically no tool support for these questions and research results are just beginning to emerge [40], [50], [12]. Based on these results, we extend and implement the handling of a range of privacy goals in the tools OFMC and AIF that we use in the evaluation of WP1.2. While this is focusing on the means necessary for verifying privacy-enhancing identity protocols, we also investigate the relationship to k-anonymity [90] (which, in a database, suggests the suppression and generalization (obfuscation) of quasi-identifiers to make an individual's data entry indistin- guishable from others) and its derivatives [78], [77], and differential privacy [54].

Task 24.2.2: Formal Semantics for Claims Languages. This task develops a formal semantics for the claim-language following the example of the CARL language [35]. The idea is to describe the amount of information specified as disclosure to the verifier by a first-order formula and establish what can be derived from several related or unrelated identity proofs. This formalizes two fundamental requirements:

I: The soundness of the implementation, that is, users cannot prove properties about themselves that do not actually hold true.

II: The privacy or completeness, that is, the server does not learn more information than users agreed to prove about themselves. In particular, a server should be unable to infer whether two identity proofs where made by the same user or not, unless users deliberately link their actions.

3.3 Research on Privacy-Friendly Audit and Data-Handling Mechanisms (Task 24.3)

Users reveal data to service providers for authentication purposes. Data-handling mechanisms determine how this data is managed by the service provider. Additionally, audits are required to ensure that service providers fulfill the pertinent data-handling policies. Task 24.3 conducts research on both topics.

Data-Handling Mechanisms. Whenever a user reveals personal data to a service provider, this data can be considered to become a resource on its own for which the user can specify

his preferred data handling policies. The service provider is then restricted in the usage of this collected personal data, and can e.g. only share such data with third parties if the user explicitly allowed this in his policy [92]. Current solutions assume that the service provider can be trusted to respect such received data handling policies. We will survey how that trust assumption can be minimized by using cryptographic mechanisms to enforce the compliance with the policies. This will be complemented by investigating on mechanisms that allow the cryptographic detection of policy violations.

Audits. In a similar vein, an identity/service provider might have to reveal logs of transactions or received presentation tokens to an external inspector for the verification or re-validation of its trustworthiness and compliance. While the authenticity of the data must be guaranteed, it is also desirable to reveal only the amount of user-specific information that is minimally required by the inspector, in order to protect the personal data of the users. For some signature schemes, mechanisms to sanitize a message without invalidating the corresponding signature already exist [10], [14], [28]. We will investigate how current privacy-enhanced credentials and eID solutions can be extended to allow for such legitimate post-processing as well.

3.4 Research on Privacy-Friendly Revocation Mechanisms (Task 24.4)

An important aspect in the trust reputation of identity providers is their ability to react on changes, in particular to revoke credentials in case they get lost or corrupted, or a user lost his right to possess a certain credential. Such revocation is more challenging when advanced identity schemes that support pseudonymous authentication are used, since therein different transactions of the same user are supposed to be unlinkable. There exist already a variety of cryptographic approaches that solves revocation in a privacy-friendly manner [26], [34], [31], [32], [83] but some important challenges remain, such as the (partial) revocation of pseudonyms, or solutions that are non-linear in the number of (revoked) users.

Task 24.4 has the following goals:

I: Investigating on new cryptographic protocols that address the above-mentioned challenges.

II: Push forward a unified framework and infrastructure that allows to address the revocation of public-key credentials in a common, technology-independent way.

3.5 Methods for Usable Privacy (Task 24.5)

This task aims at establishing methods and a corresponding framework for Usable Privacy of identity protocols in FutureID. The most important privacy requirements in this space relate to the ability of citizens to control "whether, when and to whom" [88] their personal information is disclosed, as well as their awareness and express consent on any disclosure. Such control is likely to differ amongst citizens in accordance with their social attitudes. These attitudes are determined by different cultures, norms and laws which are applicable to each citizen. The

well-established "user-centric" paradigm, analysed in [21] and realized in [36] and other identity systems, approaches this need by placing the needs of the users at the centre of the system design and empowering them with the control decisions regarding their personal information. To express these control decisions, eID systems must be usable [84] and therefore satisfy qualitative properties including the following ones:

- Learnability, efficiency, memorability, low error rates and high satisfaction [84].
- Control [85], while considering limiting factors such as information, time, and psychological deviations [8].

These limitations dictate that citizens may not be able to fully understand the different threats and risks [9] to their privacy, and be fully aware of the implications of their privacy decisions. As such, citizens may require support in making their decisions which addresses their particular model. This task takes the factors of qualitative usability, mental model and constraints into account vis a vis of the user's decision space in FutureID protocols. The decision space and the evaluation of its implications are particularly complex for composed protocols, such as in the case of eID with identity federation, and for anonymous credential systems with versatile selective disclosure. This task is therefore to establish methods and a framework to support their control decisions with respect to these limitations and ensure decisions regarding control are indeed consistent with their attitudes.

4 Extending Languages and Tools for Compositional Reasoning

Task 24.1 pursues compositional reasoning over protocols and systems. Real-world identity management systems and protocols are complex, being composed from many small components and building blocks. It is usually not feasible to reason about such complex systems as a whole, neither for the human mind nor for automated tools. We thus often need to analyze the security (and privacy) properties of components separately and in some way infer properties of the composed system. Compositional reasoning tries to establish general paradigms that allow this inference from properties of components to properties of the composite. Usually these have the form of *sufficient conditions* on the components that are relatively easy to check statically, i.e., without considering concrete runs of the system.

Besides the complexity there are two other reasons why compositional reasoning is interesting. The first is *modularity*: it allows us to *exchange* one component by another one that provides the same functionality without repeating the verification of the entire system. The second is *generality*: there may be a large or even infinite number of ways to compose components of a given library, and obviously we cannot check them all individually. Compositional reasoning can however establish such results by proving that every composition step preserved certain properties.

4.1 Three Types of Protocol Composition

Let us first review some existing work on compositional reasoning. In cryptography, composability was pursued by multiple strands of research, such as Reactive Simulatability [86] or Universal Composability [38]. In this line of work, one tries to establish composability results for components that guarantee that the component can be composed with an arbitrary environment and still guarantee its security properties.

In this task, we however focus on compositionality works in the formal methods community. Here we abstract from cryptography, treating encryptions as black boxes the intruder cannot access (except when knowing the corresponding keys). This and the focus on simple state-based or trace-based safety properties allows us to give compositional reasoning with less restrictive assumptions and without the burden to reason about indistinguishability of systems.

We speak in the following about *protocol composition* where we use protocol in the broadest sense as communicating distributed systems. There are three kinds of protocol composition that have been considered previously: parallel, sequential, and vertical composition.

Parallel composition of security protocols means that several protocols are executed over the same communication medium, where the protocols are possibly using the same key infrastructures [48, 44]. Though the protocols may be secure in isolation, their use in parallel may allow for attacks: For instance, an intruder could re-use messages or message parts from one protocol in another, if the protocols have similar message formats. In a nutshell, to achieve parallel compositionality we thus want that messages of different meanings should not be con-

fusable, i.e., we want to have disjoint message formats to have a simple syntactic check on the involved messages.

Sequential composition of security protocols means that an output of one protocol is used as input for another. This happens, for instance, when the session key output of a key exchange is used as input for a secure channel. Datta et. al. [48] have analyzed the sequential and parallel composition of protocols by composing processes, yet without considering keys. Ciobâcă and Cortier [41] analyze such compositions under consideration of shared data and employ disjointness principles introduced by Guttman and Thayer [66] to obtain a generic theorem for parallel and sequential composition.

Vertical composition was first considered by Mödersheim and Viganò [81]. The idea is that some protocols like TLS [52] establish a new *communication channel* with certain security properties, e.g., authentication of the server and confidentiality of messages. The channel means that we can basically run any *application protocol* over this channel, e.g., a login or a single-sign-on protocol. We want to verify the channel protocol independent of the concrete application protocol, and vice-versa verify the application protocol under the abstract assumption of properties of the channel, independent of its realization. This first result was extended by Groß and Mödersheim [62] to allow several layers of vertical composition and giving concrete syntactic criteria as sufficient conditions. These have again, at their core, the disjointness of message formats whenever messages have different meanings.

In FutureID we have all three kinds of protocol composition: First, key exchange, secure channel and identity management protocols are all executed over the same wire, which constitutes a parallel composition. Second, protocols are executed one after another, e.g., when a root authentication with an eID card leads to an identity federation protocol run with SAML. This constitutes a sequential composition. Third, protocols are executed on top of each other in many cases, such as when a chip authentication is executed through a secure channel established with a PACE key exchange protocol. Here we have a vertical composition.

4.2 Present Work

The question of disjointness of message formats is central to most compositional reasoning results and therefore the project has begun a deeper investigation of disjointness. While still in a early stage, first results have been achieved which we summarize in the following. We connect the abstract message terms used in the formal methods community with the actual messages used in real protocol implementations. This involves in particular reasoning about the lengths of different fields. This is an aspect that has been completely ignored in the abstract models so far. In particular we consider *formats* that consider of list of *fields* which may either be:

- Constants/tags to indicate the meaning of the message or that structure it. In particular we consider messages structured by XML-tags.
- Fields of fixed size (e.g., 64-byte random number)

- Fields of variable size that start with a length information.
- Encrypted message parts recursively have a plain-text structured according to some format.

We give a static algorithm to check whether two given formats are indeed disjoint and are uniquely parsed, i.e., there is no ambiguity in the interpretation of messages.

Based on this check we can establish a first important compositionality result: given a protocol that passes the disjointness check; if this protocol has an attack with concrete messages, then this attack can be reproduced on the more abstract model where the different formats are treated as black boxes and length fields are ignored. This bridges towards to so-called *typing results* like [69, 24, 13, 80] regarding the *formats as types*. In all these works, the main point is that we can safely move from complex realistic models to simpler abstract models: if we verify that the abstract model is free from attacks, then so is the concrete one.

4.3 Connection to APS

Here is an important connection to our Authentication Protocol Specification language (APS), designed in Task 42.3, described in the previously submitted Deliverable D42.3. The language serves

- for precisely describing a protocol;
- for deriving both a mathematical model suitable for existing verification tools; and
- for deriving an actual implementation.

Here the concept of formats is used as an important language construct for structuring plain-text messages (rather than talking about "concatenation") The abstract model works with the formats on an abstract level, the implementation instead calls a *pretty printer* for the respective format when constructing/sending a concrete message, and a *parser* when deconstructing/receiving a concrete message. We can thus say that with the first results of Task 24.1, we are able to *safely abstract from the parsing problem.*

5 Establishing Methods and Languages for Privacy Goals

Task 24.2 describes a novel approach to analyse and reason about privacy goals. We first present the context and motivation for research on new notions of privacy goals and second we explain how this innovative approach allows us to specify privacy goals in a declarative way and as a formula that fits with privacy friendly identity systems while at the same time benefitting from existing research results and verification methods.

5.1 Context and Motivation

Several formal notions of privacy have been proposed over the last decade, e.g., [4, 5, 23, 45, 51, 55, 72, 91]. Although these notions are quite different, it can be agreed that defining privacy is quite subtle and not as easy as it is supposed to be. One of the main reasons is that classical secrecy notions do not apply to data that are not secrets themselves, e.g., a vote is not itself a secret value such as a private key. Rather, the information we would like to protect is the *relation* between the (usually non-secret) values, such as for instance which voter has cast what vote.

For this reason, the vast majority of the popular approaches to formalizing privacy is based not on the question of what the intruder can deduce from a set of known messages, but rather whether he can *distinguish* two different worlds.[1] An interesting follow-up question is therefore: what is the "right" set of distinguishability questions to define privacy? For instance, in a voting protocol where each user can just vote yes or no, we may check that the intruder cannot distinguish the world where a given voter voted yes from the one where this voter voted no. However, this is not enough: even if the intruder cannot determine the votes, he should also not be able to tell whether two voters have voted the same.

When we look at privacy-friendly identity management, we have even more different kinds of data and possible relations between them, such as date of birth, home address, or different uses of the same credentials. Therefore a research problem is being confident that a given set of distinguishability questions is sufficient for privacy, that is, that we have not overlooked some possible connection the intruder could make that we would prefer him not to be able to make.

5.2 Contributions

In a first paper, we take a step back and approach the problem from a different angle. Our main goal is to find a formal description that reflects the idea of privacy in a "natural" and less technical way and that can then be related to the existing multi-dimensional notions of privacy while supporting or criticizing them. In fact, ultimately we want to use the existing results in this field, but we take the scientific liberty to first think in a slightly different direction.

[1] This is not unlike the earlier paradigm shift in cryptographic definitions from deducibility questions (such as: can the intruder obtain the plaintext of an encrypted message?) to distinguishability questions (such as: can the intruder distinguish the encryption of different chosen values?).

More specifically, we introduce a novel, simple and declarative approach to specify privacy goals, called α-β *privacy*, which is based on specifying two formulae α and β in First-Order Logic with Herbrand Universes [70].

α formalizes the intentionally released information, that is the information that we can legitimately give to the intruder, which we also refer to as *payload*. For instance, in a privacy-friendly zero-knowledge credential system (such as IBM's *Idemix* [72]) a user may prove that she is a female older than 18 years (according to an electronic passport she owns), without releasing any more information, such as her name or the precise date of birth. Hence, we have an immediate specification of the data that the user deliberately released, that is, the statement proved in the zero-knowledge proof, and it is intuitive that we then have a violation of privacy whenever the server who verified the zero-knowledge proof can derive more about the user than the user deliberately released by the proof. We must however exclude everything that is already entailed by the proved statement from this definition. For instance, the fact that the user is also over 15 years old is entailed by the proved statement and is not a violation of privacy, but if the server is able to derive that the user is actually over 21 years, then there is a violation. It is thus quite natural to formalize such statements as formulae in some logic and to define privacy as the inability of the intruder to derive statements that are not entailed by what the users have released.

As a counterpart to the "ideal knowledge" provided by the payload α, we also need the *technical information* β, which represents the "actual knowledge" that the intruder has, describing the information (such as names, keys) that he initially knows, which actual cryptographic messages he has observed and what he knows about these messages. For instance, he may be unable to decrypt a message but anyway know that it has a certain format and contains certain (protected) information, such as a vote.

α-β privacy then means that the intruder cannot derive any "non-technical" statement from β that he cannot derive from α already. We believe that this is indeed a simple way to define privacy, and is a more declarative way to talk about privacy than distinguishability of frames. Essentially, the modeler should not think about what technical information the intruder could exploit, but rather what information he is fine to release (α) and what messages are actually exchanged (β).

Another interesting and very declarative feature of our approach is that it is straightforward to model what happens when two intruders collaborate and share their knowledge. α-β privacy allows us to formalize this simply by taking the logical conjunction of the formulae describing the knowledge that the two intruders have. Reflecting in a natural way the requirements of the systems are that although the best technology cannot prevent dishonest agents from pooling all the information that they were intentionally given and deriving all possible conclusions from that, we can ask that they cannot derive more than that.

We describe by a variety of examples how α-β-privacy can be used in practice, and define transition systems based on it. Even though α-β privacy does not directly contain a notion of distinguishing between worlds, there is a close relationship to static equivalence of frames that we investigate formally. This allows us to justify (and criticize) the specifications that are currently used in verification tools and obtain partial tool support for α-β privacy. We also

prove several results that help in reasoning about α-β privacy in general and give a decision procedure for a fragment of it.

6 Research on Privacy-Friendly Audit and Data-Handling Mechanisms

Task 24.3 conducts research on data-handling mechanisms and on audits. Data-handling mechanisms determine how user data is managed by the service provider. Additionally, audits are required to ensure that service providers fulfill the pertinent data-handling policies.

6.1 Data Handling Mechanisms

To facilitate the design of data-handling policies and mechanisms, privacy-preserving authentication and access control should be employed in order to minimize the data that data-handling policies and mechanisms must deal with. Therefore, we have designed several protocols for privacy-preserving authentication and access control:

Oblivious transfer with access control. When users access a data record from a provider, it is important that users have their privacy protected, i.e., providers must not become privy to which user accesses which data record, and providers need to be ensured that only users with the proper attributes are allowed access. A number of cryptographic protocols have already been proposed that address this problem. They all meet the first requirement by employing an adaptive oblivious transfer (OT) protocol and, to fulfill the second requirement, extend it with an attribute-based access control enabling users to prove that they satisfy the policy for the item they want to (obliviously) access. To this end, users are issued anonymous credentials attesting them their attributes or roles during setup by either the provider or a dedicated issuer. We call such schemes oblivious transfer with access control (ACOT).

More concretely, the security properties offered by ACOT can be described as follows:

User Privacy: The database cannot tell which user makes a query, nor can it tell which record is being accessed. That is, the database only learns that some user accesses some record for which the user priorly obtained the necessary credentials. If the database colludes with the issuer and potentially with other users, then they can only try to identify the user or her selection based on which credentials were issued to whom, and which credentials are necessary to successfully access which record.

Database Security: A cheating user alone cannot access a record for which she does not have the necessary credentials. Colluding users cannot pool their credentials, meaning that they cannot access any records that none of them would have been able to obtain individually. If the issuer colludes with one or more users, they can only obtain as many records from the database as the number of transfer queries that were performed.

We provide the first ACOT protocol [7] that is provably secure in the universal composability (UC) framework and hence offers superior security guarantees. Our scheme assumes a CRS and is secure under the XDLIN and SXDH assumptions (which are simple extensions

of the DDH assumption). Note that using a CRS is the best we can hope for as it is known that UC-secure adaptive OT protocols are impossible without any setup assumptions.

We additionally provide the most efficient UC-secure adaptive OT (UC-AOT) scheme of those that are secure under the standard non-q-type assumptions. We prove it secure under two variants of the DDH assumption. Our UC-AOT scheme employs a new structure-preserving signature scheme, which can be considered a contribution of independent interest. Finally, we point out some oversights in prior security definitions for adaptive OT in the UC model and discuss the impossibility of dynamic adversaries as well as extensions that allow the adversary to abort a transfer or the provider to add records iteratively.

Concepts and Languages for privacy-preserving authentication. Privacy-preserving authentication mechanisms are based on advanced cryptographic primitives such as anonymous credentials, minimal disclosure tokens, self-blindable credentials, or group signatures. One possible reason for the slow adoption of privacy-preserving authentication technologies might be that the various schemes described in the literature have a large set of features where similar features are often called differently or are realized with different cryptographic mechanisms. Many of the features such as credential revocation, efficient attribute encoding, or anonymity lifting even require a combination of separate cryptographic protocols. This makes these technologies hard to understand and compare and, most importantly, very difficult to use.

We provide unified definitions of the concepts and features of the different privacy-preserving authentication mechanisms. We will refer to this unification as privacy-preserving attribute-based credentials or Privacy-ABCs. Our definitions abstract away from the concrete cryptographic realizations but are carefully crafted so that they can be instantiated with the different cryptographic protocols—or a combination of them. To enable the use and integration of Privacy-ABCs in authentication and authorization systems, we further present cryptography-agnostic definitions of all concepts as well as a language framework with data formats for, e.g., policies and claims. All languages are specified in XML schema and separate the abstract functionality expected from the underlying cryptographic mechanisms from the opaque containers for the cryptographic data itself. Thus, these languages allow application developers to employ Privacy-ABCs without having to think about their cryptographic realization, similarly to how digital signatures or encryption can be used today.

Fair Private Set Intersection. Private set intersection (PSI) enables two mutually untrusted parties to compute jointly the intersection of their private input sets. The majority of the protocol proposed to solve the PSI problem are single-output protocols, i.e. only one party obtains the intersection and the other party gets nothing. However, there are many motivating scenarios in which both parties want to know the intersection, e.g., two real estate companies would like to identify customers (e.g., homeowners) who are double-dealing.

A mutual PSI protocol must be fair, i.e. if one party knows the intersection, the other party should also know it. To efficiently achieve fairness, most fair cryptographic protocols are optimistic which requires help from an off-line arbiter who is a trusted third party. The

arbiter only participates if one party unfairly aborts the protocol and can recover the output from the protocol for the honest party. Incorporating optimistic fairness in PSI protocols is not easy for two reasons: firstly, although there is a generic structure, there is no generic construction for optimistic fair protocols. Secondly, the arbiter usually has to get access to some private information and therefore has to be fully trusted. However, in reality it is hard to find such a fully trusted third party.

We present the first fair mutual PSI protocol. The protocol has built-in support for optimistic fairness and does not require setup assumptions such as certified input sets. In addition, the third party acting as the arbiter can resolve disputes without knowing the private inputs or the output of the PSI protocol. Hence we can significantly reduce the trust placed on the arbiter.

Protocols for UC-secure circuit evaluation. Privacy-preserving authentication and access control protocols often consist of several building blocks. If security proofs of the building blocks consider only a single instance, all security guarantees are lost if it is run concurrently with other protocols or with itself. Better security guarantees can be obtained when using composability frameworks, which ensure that protocols proved secure in the framework remain secure under arbitrary composition. This also simplifies the design of protocols: high-level protocols can be composed from building block protocols and the security proofs of the high-level protocols can be based on the security of the building blocks. Unfortunately, protocols proven secure in such composability frameworks are typically an order of magnitude less efficient than their traditional counterparts with "single-instance" security. Moreover, most UC-secure schemes and protocols found in the literature can not be used as building blocks for higher-level protocols because they do not offer the proper interfaces.

We provide practically useful UC-secure building block protocols that provide interfaces so that parties in higher-level protocols can prove to each other that their inputs to one building block protocol correspond to the outputs of another building block protocol. More precisely, we provide a set of two-party protocols for evaluating an arithmetic circuit with reactive inputs and outputs. Our protocols evaluate an arithmetic circuit modulo a composite number n, where n is a product of two large safe primes that is assumed to be generated by a trusted third party, and whose factorization remains otherwise unknown. Our protocols are more efficient than existing UC-secure two-party circuit evaluation protocols which were designed to be used as stand-alone protocols.

Optimally Private Access Control. Anonymous credential schemes allow Alice to prove to Bob that her credentials fulfill the access rule of a resource in a privacy-friendly manner. Particularly, Alice proves possession of credentials without revealing the credential's attributes, yet proving that the attributes fulfill the constraints described in the access rule. Additionally, Bob cannot link different proofs of possession of the same credential even if he colludes with the credential issuer, which allows the implementation of anonymous or pseudonymous access control. However, existing access control protocols based on anonymous credentials can be improved in terms of privacy and efficiency.

Current anonymous credential schemes typically reveal extra information, such as the identity of the issuers of Alice's credentials and the access rule employed. Such informa-

tion reduces the anonymity set, which could eventually lead to Alice's identification. In addition, the efficiency of selective attribute disclosure is not optimal.

We propose the first efficient credential-based access control protocol that is optimally private, i.e., the service provider only learns whether a user should be granted access or not. Optimal privacy requires, on the one hand, to hide all the information about the credentials whose possession is proven, such as the issuer's identity, the credential type and the number and value of the credential's attributes. On the other hand, it requires to hide the access rule used to prove access rights, which implies that neither information about the number and type of credentials being shown, nor about the number and type of attribute constraints being enforced, should be revealed. Additionally, our protocol provides efficient attribute disclosure for all attributes.

Existing eID Solutions. Within the FutureID project as use-cases we investigated (parts of) the Austrian eID ecosystem [95] and the architecture proposed in the Secure Identity Across Borders Linked (STORK) project [96][2]. In particular, we investigated how data minimization techniques could be integrated into such infrastructures without significant changes in the infrastructure and at the same time moving components which suffer from scalability issues into the public cloud (which is assumed to operate honest but curious). It turns out that a pragmatic solution, besides the use of anonymous credentials, seems to be to apply redactable signature schemes [73, 89] in combination with suitable proxy re-encryption primitives [15, 60]. In doing so, service provider only receive information required for service provisioning (data minimization) and nothing beyond. However, it is an open question how to cryptographically enforce that a service provider does not pass received (identity) information to any other (untrusted) party.

6.2 Audits

Homomorphic or malleable signature schemes [73, 11, 25, 49] are signature schemes which support the evaluation of certain classes of functions on authenticated data while preserving authenticity. Well known classes are redactable signatures [73, 89] and sanitizable signatures [14, 27]. The former supports removal of message parts by any third party while the latter allows replacements of designated parts by a designated party (being in possession of a corresponing trapdoor). Consequently, in these schemes, given a message and corresponding signature, one can compute a signature for a submessage (or allowed modified message) derived from the original message without having access to the private signing key. Consequently, to some extent, a signer can control how signed data can be modified prior to passing this data to some other (untrusted) third party.

Within the FutureID project we proposed two somewhat related signature schemes, namely blank digital signatures (BDS) [67] and warrant-hiding proxy signatures (WHPS) [68]. WHPS basically allow to delegate the signing rights for a set of messages to a proxy as in conventional proxy signature schemes. The proxy then chooses one message out of this set and issues a signature on it. Upon verification anyone is able to verify the validity of such a signature whilst

[2] `https://www.eid-stork.eu/`

not learning anything about the remaining message space. BDS allow for the delegation of the signing rights for a so called template, constituting of fixed and exchangeable elements, i.e., a set of choices, to a proxy. The proxy is then able to sign any instantiation of such a template which corresponds to the template, i.e., contains all fixed elements and a single choice for each exchangeable element, on behalf of the originator. Upon verification, anyone is able to verify the validity of the signature whilst not learning anything about the unused choices in the exchangeable elements. Interestingly, it seems that both signature schemes can be obtained from anonymous credentials, whilst the opposite direction is not true. However, some recent works show that there is also a connection between malleable signature schemes and anonymous credentials [39, 37], whilst one typically requires specific properties of the malleable signature scheme to hold - in particular unlinkability.

Furthermore, quite recently other interesting approaches have been proposed. In [20], Bellare and Fuchsbauer propose an approach which basically allows for defining policies enforcing certain properties on signed messages. In order to issue correct signatures w.r.t. a policy, one additionally is required to be in possession of a signing key for this particular policy. Their general construction allows for defining policies for any language in NP (the set of decision problems where the "yes"-instances can be accepted in polynomial time by a non-deterministic Turing machine), i.e., it can be verified in polynomial time whether a particular message corresponds to a certain policy, which is then restricted to group dependent languages due to using Groth-Sahai proofs [63] in their instantiation. Somewhat related, Backes et al. [18] propose a model for delegating the signing rights for messages being derivable from an initial message by applying a function to the message. It is interesting to study such signature schemes, their expressiveness, their relation and application to anonymous credentials as well as their application to problems within the FutureID project in the future.

On the other hand, we have designed a mechanism for privacy-preserving audits that does not employ redactable or sanitizable signatures. In our audit mechanism, from all the attributes that the user reveals to the verifier, the user determines the attributes that the verifier is able to reveal to the auditor. Moreover, from all the attributes that the verifier can reveal to the auditor, the verifier determines the subset that is revealed to the auditor.

Our audit mechanism works as follows. First, a user obtains from an issuer an attribute-based credential that certifies L attributes. When the user needs to prove possession of some attributes to a verifier, the user computes a presentation token and sends it to the verifier.

The presentation token is computed as follows. Let $D \subseteq [1, L]$ be the subset of attributes that the user reveals to the verifier and that the verifier is able to transfer to the auditor. Let $F \subseteq [1, L]$ be the subset of attributes that the user reveals to the verifier but that the verifier cannot transfer to the auditor. The user commits to attributes in $D \cup F$ and computes a Fiat-Shamir zero-knowledge proof that the committed attributes equal those in the credential. For the attributes in D, the user reveals to the verifier the openings of their respective commitments. For the attributes in F, the user computes Fiat-Shamir zero-knowledge OR-proofs (i.e, a proof that involves a disjunction) that the user knows either a valid commitment opening or the verifier's secret key. These proofs convince the verifier of the validity of the revealed attributes in F but prevent the verifier from disclosing those attributes to the auditor in a verifiable manner.

In the audit phase, the verifier reveals to the auditor the commitments to the attributes in $D \cup F$ and the proof that the committed attributes were certified by the issuer. For the attributes in D, the verifier chooses the ones that he wishes to reveal to the auditor and sends the auditor the openings of the commitments. For the attributes in F, the verifier cannot convince the auditor that he knows those attributes because the OR-proofs could have been produced by the verifier.

Finally, we also extend the Identity Mixer anonymous credential system to enable signatures on committed graphs and zero-knowledge proofs of knowledge of those signatures. This is relevant to FutureID in the area of audits: The method allows certification of an infrastructure's topology (such as a large-scale identity management system) and proofs of its properties, without disclosing the entire graph to the verifier.

7 Research on Privacy-Friendly Revocation Mechanisms

Following [32], there are many reasons why a credential needs to be revoked. The user might have lost her right to carry the credential, the secret key underlying the credential might have been compromised, or just because the attributes stated in the credential became outdated.

A possible solution to revocation in the case of non-anonymous credentials is to blacklist all serial numbers of revoked credentials in a so-called certificate revocation list. In such a setting users have then to demonstrate that their credentials are not contained within this blacklist. Similarly, one could follow a whitelist approach, i.e., maintaining a list containing valid credentials, and users have to demonstrate that their credential is contained within this whitelist. In non-anonymous approaches using revocation lists containing serial numbers, as for instance used within PKIX (a public key infrastructure that use X.509 certificates), this is a simple lookup for a verifier when a credential is shown. However, this does not work for anonymous credentials as showing a serial number would break anonymity and unlinkability.

A solution inspired by revocation lists is the use of so-called dynamic accumulators [31, 34]. Here, all valid serial numbers are accumulated, i.e., compressed, into a single value that is then published. In addition, dynamic accumulators provide a mechanism that allows to dynamaically add and remove users from such an accumulator. Furthermore, a user can prove that the serial number of her credential is contained (or not contained) in the accumulated value. Whenever a credential is revoked, a new (updated) accumulator value is published that no longer contains the revoked serial number (this can be done analogously for the whitelist approach).

Accumulator based schemes require, however, that users keep track of the changes to the accumulator and their so called witnesses to be able to execute their validity proofs. Another drawback is that they all make proving and verifying ownership of credentials less efficient (typically about a factor of 2 or worse), as not only possession of the credential has to be proven but also that it is still valid w.r.t. the revocation list/accumulator.

Future work within the FutureID project includes investigating efficient solutions to credential revocation.

8 Methods for Usable Privacy

Task 24.5 focuses on establishing new methods for usable privacy. In Section 8.1, we review the findings and recommendations on the usability of privacy enhancing identity management systems from the PRIME and PrimeLife projects. In Section 8.2, we pursue a case study with privacy in real-world environments and in Section 8.3 we lay the foundations for new insights in usability for privacy by helping users establish appropriate mental models.

8.1 Usable privacy considerations in PRIME and PrimeLife

PRIME and PrimeLife proposed user-centric identity management systems. While PRIME made usable privacy recommendations, PrimeLife identified challenges in addressing user privacy requirements and evaluated their proposed solution.

8.1.1 PRIME

PRIME proposed audience segregation and user control as two key aspects of user-centric identity management [87]. Audience Segregation: Individuals have different identities that are used in different settings in society such as citizen, daughter, friend or employee. While exercising these identities, they explicitly or implicitly disclose information about themselves. However, in the online environment, 'simple' partial identities can be aggregated into rich compound identities from data linked to identifiers such as names and IP-addresses, i.e., the addresses employed by the Internet Protocol. Digital personae are easily copied, merged and manipulated exposing these personae to 'audiences' that should not be able to obtain personal data and allowing use of data out of context which in turn can lead to reputation damage. Having different partial identities is therefore a social necessity and an essential aspect of informational privacy. User Control: User control ranges from what gets disclosed to whom, up to very strong positions such as the right to informational self-determination. User control in design can be decomposed into the following sub-components: information to the user, consent of the user, user access, correction, erasure, and objection, and security and trust. However, affordability and skill levels influence the ability to access and use technology that enable user control. Affordability is related to the perceived usefulness of the system, that is information about the product, and comprehensibility of its features can influence the perceived affordability of a service. Users should also be able to use an application with minimal amount of training. Moreover there are no homogeneous user group, and skills change between social groups or nations. Identity management systems should therefore be usable for non-skilled users and provide satisfactory default settings. Identity management systems must also pay attention to the contextuality of privacy – that is give individuals the possibility to change privacy settings according to context.

8.1.2 PrimeLife

The development of usable privacy-enhancing Identity Management poses several HCI (Human Computer Interaction) challenges. The PrimeLife project conducted research on methods for usable privacy and summarized the problems that need to be addressed in [57]. In particular, the typical HCI fallacies that were experienced when developing user interfaces for Privacy-Enhancing Technologies (PETs) and guidelines for the design of usable privacy-enhancing technologies are described.

The major issues that were experienced and that should be considered during user interface design and testing can be summarized as follows:

- Many users cannot differentiate whether data is stored on the user side (under the user's control) or on a remote services side.
- Many users do not understand to which network entities personal data flows during online transactions.
- Privacy warnings can cause rushed and unwanted user reactions and thus need to be designed with care.

These are mental model issues which are difficult to solve for the novel privacy-enhancing technologies concept, which is unfamiliar for the users. This is especially true for those PETs, for which no obvious real world analogies exist, such as for instance for anonymous credentials and their selective disclosure properties.

To address these issues, PrimeLife provides guidelines for the design of user interfaces for privacy-enhancing technologies as well as methods for evaluating those user interfaces.

User Interface Design. User interfaces should address the challenge to evoke the correct mental model in regard to where data are transferred to and where they are processed. PrimeLife provides HCI heuristics for PETs, which adapt, extend and exemplify the classical list of Nielsen's Usability Heuristics for the PET domain. Besides, PrimeLife provides HCI Patterns for PETs, which give best practice solutions for the PET user interface design and which should be applied in combination with the Usability Heuristics.

User Interface Evaluation. PrimeLife explains that the mediation of trustworthiness, intercultural differences and an understandable terminology to be used in UIs are challenges to be taken into consideration. When recruiting test participants, aspects such as their cultural and technical background need to be considered. After the tests, PrimeLife proposes the use of post test questionnaires and PET-USES (Privacy-Enhancing Technology Users' Self-Estimation Scale), which were developed to obtain a more accurate account of the experience and opinions of the test participants.

PrimeLife further investigated the user-centred challenges of privacy-enhanced identity management. We present PrimeLife's recommendations on tackling the challenges such as limited user

knowledge of privacy features, technologically driven development of privacy tools, understanding privacy tools related terms, wrong mental models of privacy tools, privacy as a secondary task and complex mechanisms that are hard to understand [59]. We then describe PrimeLife's findings of usable privacy evaluation.

Limited user knowledge of privacy tools: PrimeLife found that the design can aid users in understanding privacy tools and privacy concepts when they are offered with clear interfaces and structures that display privacy aspects and potential threats in an understandable way [56]. PrimeLife used the principle of designing a standard view prototype for novice and average users and an expert view for technologically minded persons who can setup and configure their preferences.

Technologically driven development of PETs: PrimeLife also found that increasing the user's knowledge about privacy eases their understanding of privacy concepts. From works on privacy patterns [58], PrimeLife proposed clear structures, visual aids such as icons that can help understanding the technology.

Understanding of privacy technology related terms: PrimeLife underlined the need for users to understand privacy related words and providing further explanation for unknown privacy terms. PrimeLife found that some terms are easiest to understand (such as privacy protection, required data, digital traces, identity management and full privacy policy) and some terms were hard to understand (such as anonymous credentials, privacy preference, linkability, privacy enhancing).

Wrong mental models of privacy tools: PrimeLife proposes a requirement gathering process to form a picture of users' understanding of how privacy and privacy tools work and argues that designs must ensure mental models are respected in user interfaces.

Privacy as a secondary task: PrimeLife highlights the need for tools that do not require much time and effort from the user side and that do not disturb users' workflow with unnecessary pop-ups for example.

Complex mechanisms are hard to understand: Users are overwhelmed by the complexity and amount of information presented – thus there is a need for designers to have a clear picture of expectations and knowledge of users and to adapt interfaces to this knowledge. To reduce complexity, PrimeLife suggests presenting only important information through self-explanatory visual concepts as icons.

Influence of mental models on users' comprehension of anonymous credentials. PrimeLife analysed the effects of users' mental models on their understanding of the data minimisation property of anonymous credentials in the context of an e-shopping application scenario [93]. Results showed that the card-based approach led to significantly more errors of addition (users believe that they disclosed more information than they actually have) whereas the attribute based approach led to more errors of omission (that is users underestimated the amount of data that they have disclosed).

Privacy policy icons. PrimeLife proposed privacy icons as a way to help users understand and read privacy policies and to improve their awareness of what is happening with their personal data [71]. PrimeLife has developed icon sets for different use cases such as e-commerce, social networks and handling of e-mails. However PrimeLife found that the use of icons and the incorporation of machine-readable policies as well as their relationship to each other and to their practice of presenting privacy policies has to be spelled out so as to avoid users' getting a totally different picture of the privacy processing than those reading the privacy policy in natural language.

8.2 Case Study: eID Privacy in SEMIRAMIS

Electronic Identity (eID) addresses electronic identity management. It consists of a highly secure end-to-end channel of communication between the user, the service provider and the identity provider allowing the user (who can use a personal, portable and secure document) to gain access to her user eID.

The use of e-Services, which entails the electronic exchange and recording of data, must be adopted with the premise of protecting the privacy and identity of the individuals. Consequently effective security, privacy and trust functionalities are key requirements and privacy preserving eID based architecture is necessary. The SEMIRAMIS project has done so through the design of its architecture and provides a fundamental starting point for preserving eID privacy.

8.2.1 Introduction to SEMIRAMIS project

SEMIRAMIS [2] is an innovative project under the Information Communication Technologies Policy Support Programme (CIP-ICT-PSP), executed by a consortium of 9 European industrial and academic organizations. SEMIRAMIS is built on mature technical solutions (such as the architectures provided by SWIFT [3] and GÉANT [1]) and standards. The goal of this SEMIRAMIS is to define an infrastructure for providing e-Services involving secure cross-border data exchange [16]. SEMIRAMIS caters for confederated identity management across three types of federations or domains: education, governmental, and telecommunications sectors in different countries. These domains need to interact with one another in order to provide composed e-Services across Europe.

8.2.2 Privacy in SEMIRAMIS

SEMIRAMIS aims to provide for users' privacy needs, ensuring an appropriate level of privacy and compliance with data protection regulations, as well as usability aspects, all of which are key to the uptake of transnational e-Services. SEMIRAMIS ensures an appropriate level of privacy first by establishing a privacy-preserving architecture with a focus on the different privacy-enhancing mechanisms and technologies used and second through the use of pseudonyms and privacy measures protecting the release of user information to third parties. SEMIRAMIS also applies end-to-end privacy based on a "need to know basis" approach, for instance for information

exchange on a need-to-know-basis each component only has access to the information and data needed to perform the functionality of the component. End-to-end encryption of attribute values [79] is also applied to avoid intermediate elements to see attribute values and know user information, which is not really needed for the purpose of providing the service and to allow the user to validate the requestor certificate.

The use of pseudonyms is convenient in the case of systems where the real identity of the user must be preserved, while, at the same time, some level of traceability of the user's past actions is required. Pseunonyms are an easy way to avoid leaking user identity information by using a user-unique identifying key at the Identity Provider that guarantees the users' anonymity during the communications. SEMIRAMIS also provides control mechanisms for information release, that is, the users define policies that control the release of their information to the requesting parties and are asked for user consent before releasing their personal information. SEMIRAMIS' architecture also provides an auditing functionality to users that guarantees that their data is only accessed by the allowed entities.

8.2.3 Considerations for FutureID

Some of the privacy preserving mechanisms in SEMIRAMIS could be applied to FutureID while others would not suit FutureID's requirements – however they might influence usability:

- Need to know basis for the information exchange. This feature should be considered in FutureID. Each component should only have access to the needed information for performing the functionality. This measure does not have an impact on the usability.
- End-to-end encryption. The use of this feature is possible in FutureID, but the implications should be analyzed as the intermediate components between the service provider and the Identity provider might need to know the contents of the interchanged messages and this would not be possible if the content is encrypted.
- This feature would help to improve the user's privacy and at the same time not have an effect on the usability.
- User Consent. This feature is mandatory due to legal restrictions (data protection laws) and it can be implemented more or less explicitly. The more explicit requirement leads to higher impact on usability.
- Auditing functionality. This feature should be implemented in FutureID in order to provide a guarantee of privacy to users. This feature does not affect usability.

8.3 Mental Models for Usable Privacy

The previous sections highlight the complexity of designing usable privacy-enhancing identity management systems. The recommendations and findings are important to the privacy research space. However, there is a lack of research in how designs affect reasoning and evaluation of

privacy online and consequent behaviour and how tapping into people's use of mental models would influence privacy elaboration and decisions.

Mental models can facilitate reasoning about privacy and identity and enable individuals to experiment with small-scale models built from available information. Mental models can also help to structure and link different privacy experiences, behaviours, perceptions and help to build cause and effect relationships. The previous sections provide support for the usefulness of mental models in usable privacy research.

Privacy enables individuals to control access to their self and to information about themselves. It thus allows individuals to form varying relationships and to exercise different roles in society. Privacy is implicit within human behaviour and allows the existence of a dialectical state that allows individuals to be both connected and autonomous. Therefore individuals' privacy requirements are dynamic in that they require different privacy needs in different contexts at different points in time and depending on what information is shared with whom and what the recipient can do with the information.

Individuals' ability to represent the world accurately and to reason about online privacy, however, is limited and the reasoning is unique to each individual. Hence mental models conceived as a cognitive structure that forms the basis of and facilitates reasoning and decision making and characterized as small-scale, incomplete representations of reality can enable individuals' unique elaboration of their online privacy and identity that otherwise would involve complex elaboration and high cognitive effort. Acknowledging individuals' limited ability to comprehend complex systems, mental models as 'internalised, mental representations of a device or idea' [75] can help individuals make sense of their online privacy. As a simpler version of reality constructed by the mind, mental models can help to anticipate events, to reason and underlie explanations [46] and to interact with the world. They are constructed by individuals based on their personal life experiences, perceptions, and understandings of the world and provide the mechanism through which new information is filtered and stored.

Each mental model represents a possibility and captures what is common to all the different ways in which the possibility may occur [74]. Mental models are based on a principle of truth representing only those situations that are possible, and each model of a possibility represents only what is true in that possibility according to the proposition. Therefore individuals can infer that a conclusion is valid if it holds in all the possibilities and premises. Reasoning about privacy with mental models would therefore also help individuals to refute invalid inferences.

Mental models are context dependent and highly dynamic and allow individuals to adapt to continually changing circumstances and to evolve over time through learning. This process is akin to how individuals assess and elaborate about their privacy over time and how identity is constructed. Moreover, online designs can affect how mental models of privacy are depicted and which aspects of individuals' privacy mental model become more prominent and influence privacy decisions.

Mental models have a structure that corresponds to the structure that they represent for instance, mental models could enable a hierarchical way of looking at the privacy of one's identity and partial identities. They also have a content formed through a process of 'gluing together' of

or experimenting with specific behaviours in the environment or of a system. Mental models can hence help individuals to bring together different pieces of information to help privacy decision such as who has access to personal information, when and how, who owns the personal information, the context of disclosure and privacy means available. With each application or run of the privacy mental model, relations are formed between known and unknown attributes/information in the design and cause and effect relationships are created and tested. Therefore the structure of the mental model will be constrained by perceptions of privacy and prior knowledge, (that is previous privacy experience will shape one's privacy requirements) but tools of thought such as analogies and metaphors can function as help to structure unfamiliar domains with regards to privacy.

A privacy mental model framework or theory will therefore be valuable for usable privacy design in FutureID since it taps into the processes people use to reason and make decisions. As small-scale representation, mental models will not overwhelm users. It will cater for varying user skillset and ability and will alleviate the high cognitive effort required to understand complex privacy tools and solutions. Mental model consideration in the design will influence perception of control of privacy and privacy evaluation, what information is seen relevant to privacy decisions, highlight privacy tools, and will influence how users' privacy evolves. The framework will also support the aspects of contextuality of disclosed personal information, privacy and identity, the perceived usefulness of tools, and comprehensibility of features. Therefore identification and evaluation of mental models of privacy and identity will be valuable to usable and effective privacy designs for FutureID. The audience segregation and user control concepts of PRIME, the findings of PrimeLife (such as the ones on privacy icons) and the recommendations of SEMIRAMIS (such as pseudonyms as representation of identity roles and the auditing functionality) will be valuable to further research on mental models of privacy and identity.

9 Conclusion

This deliverable describes the research conducted during the first year of the FutureID project in the five tasks described in work package 24: extending languages and tools for compositional reasoning (Task 24.1), establishing methods and languages for privacy goals (Task 24.2), development of privacy-friendly audit and data-handling mechanisms (Task 24.3), development of privacy-friendly revocation mechanisms (Task 24.4) and development of methods for usable privacy (Task 24.5).

In Task 24.1, we have shown that at the core of many compositionality questions is the disjointess of message formats and we have first results in a draft paper status. In Task 24.2, we have introduced a new notion of privacy goals that links the classical trace based with the static equivalence based approaches. In Task 24.3, we provide a bundle of designs for privacy-preserving audit and data-handling mechanisms. In Task 24.4, work on revocation mechanisms is ongoing. Finally, in Task 24.5, we have analyzed existing work and we pursue a case study with privacy in real-world environments and lay the foundations for new insights in usability for privacy by helping users establish appropriate mental models.

10 Abstracts of Research Papers

10.1 Extending Languages and Tools for Compositional Reasoning (Task 24.1)

We do not have publications for this task so far.

10.2 Establishing Methods and Languages for Privacy Goals (Task 24.2)

In this task, we have so far one publication:

> Sebastian Mödersheim (DTU), Thomas Groß (UNEW), and Luca Viganò. *Defining Privacy is Supposed to be Easy*, at LPAR 2013, to appear.

Abstract. Formally specifying privacy goals is not trivial. The most widely used approach in formal methods is based on the static equivalence of frames in the applied *pi*-calculus[3], basically asking whether or not the intruder is able to distinguish two given worlds. A subtle question is how we can be sure that we have specified all pairs of worlds to properly reflect our intuitive privacy goal. To address this problem, we introduce in this paper a novel and declarative way to specify privacy goals, called α-β privacy, and relate it to static equivalence. This new approach is based on specifying two formulae α and β in first-order logic with Herbrand universes, where α reflects the intentionally released information and β includes the actual cryptographic ("technical") messages the intruder can see. Then α-β privacy means that the intruder cannot derive any "non-technical" statement from β that he cannot derive from α already. We describe by a variety of examples how this notion can be used in practice. Even though α-β privacy does not directly contain a notion of distinguishing between worlds, there is a close relationship to static equivalence of frames that we investigate formally. This allows us to justify (and criticize) the specifications that are currently used in verification tools, and obtain partial tool support for α-β privacy.

10.3 Research on Privacy-Friendly Audit and Data-Handling Mechanisms (Task 24.3)

- Christian Hanser and Daniel Slamanig, "Blank Digital Signatures," in *Proc. of 8th ACM Symposium on Information, Computer and Communications Security (ASIACCS 2013)* [67].
  Abstract. In this paper we present a novel type of digital signatures, which we call *blank digital signatures*. The basic idea behind this scheme is that an originator can define and sign a *message template*, describing fixed parts of a message as well as multiple choices for exchangeable parts of a message. One may think of a form with blank fields, where for such fields the originator specifies all the allowed strings to choose from. Then, a proxy is

[3] *pi*-calculus is a process calculus, i.e., a tool for the high-level description of interactions, communications, and synchronizations between a collection of independent agents or processes.

given the power to sign an *instantiation* of the template signed by the originator by using some secret information. By an instantiation, the proxy commits to one allowed choice per blank field in the template. The resulting *message signature* can be publicly verified under the originator's and the proxy's signature verification keys. Thereby, no verifying party except the originator and the proxy learn anything about the "unused" choices from the message template given a message signature. Consequently, the template is hidden from verifiers.

We discuss several applications, provide a formal definition of *blank digital signature schemes* and introduce a security model. Furthermore, we provide an efficient construction of such a blank digital signature scheme from any secure digital signature scheme, pairing-friendly elliptic curves and polynomial commitments, which we prove secure in our model. We also provide a detailed efficiency analysis of our proposed construction supporting its practicality. Finally, we outline several open issues and extensions for future work.

- Christian Hanser and Daniel Slamanig, "Warrant-Hiding Delegation-by-Certificate Proxy Signature Schemes," in *Proc. of 14th International Conference on Cryptology in India (INDOCRYPT 2013)* [68].
  Abstract. Proxy signatures allow an entity (the delegator) to delegate his signing capabilities to other entities (called proxies), who can then produce signatures on behalf of the delegator. Typically, a delegator may not want to give a proxy the power to sign *any* message on his behalf, but only messages from a well defined *message space*. Therefore, the so called *delegation by warrant* approach has been introduced. Here, a *warrant* is included into the delegator's signature (the so called certificate) to describe the message space from which a proxy is allowed to choose messages to produce valid signatures for. Interestingly, in all previously known constructions of proxy signatures following this approach, the warrant is made explicit and, thus, is an input to the verification algorithm of a proxy signature. This means, that a verifier learns the entire message space for which the proxy has been given the signing power. However, it may be desirable to hide the remaining messages in the allowed message space from a verifier. This scenario has never been investigated in context of proxy signatures, but seems to be interesting for practical applications. In this paper, we resolve this issue by introducing so called *warrant-hiding proxy signatures*. We provide a formal security definition of such schemes by augmenting the well established security model for proxy signatures by Boldyreva et al. Furthermore, we discuss strategies how to realize this warrant-hiding property and we also provide two concrete instantiations of such a scheme. They enjoy different advantages, but are both entirely practical. Moreover, we prove them secure with respect to the augmented security model.

- Bernd Zwattendorfer and Daniel Slamanig, "On Privacy-Preserving Ways to Porting the Austrian eID System to the Public Cloud," in *Proc. of 28th IFIP TC-11 International Information Security and Privacy Conference (SEC 2013)* [95].
  Abstract. Secure authentication and unique identification of Austrian citizens are the main functions of the Austrian eID system. To facilitate the adoption of this eID system at online applications, the open source module MOA-ID[4] has been developed, which manages

[4]MOA-ID is a module for online application used to uniquely identify and authenticate users securely who

identification and authentication based on the Austrian citizen card (the official Austrian eID) for service providers. Currently, the Austrian eID system treats MOA-ID as a trusted entity, which is locally deployed in every service provider's domain. While this model has indeed some benefits, in some situations a centralized deployment approach of MOA-ID may be preferable. In this paper, we therefore propose a centralized deployment approach of MOA-ID in the public cloud. However, the move of a trusted service into the public cloud brings up new obstacles since the cloud can not be considered trustworthy. We encounter these obstacles by introducing and evaluating three distinct approaches, thereby retaining the workflow of the current Austrian eID system and preserving citizens' privacy when assuming that MOA-ID acts honest but curious.

- Bernd Zwattendorfer and Daniel Slamanig, "Privacy-Preserving Realization of the STORK Framework in the Public Cloud," in *Proc. of 10th International Conference on Security and Cryptography (SECRYPT 2013)* [96].
  Abstract. The STORK framework – enabling secure eID federation across European countries – will be the dominant identification and authentication framework across Europe in the future. While still in its start up phase, adoption of the STORK framework is continuously increasing and high loads can be expected, since, theoretically, the entire population of the European Union will be able to run authentications through this framework. This can easily lead to scalability issues, especially for the proxy-based (PEPS) approach in STORK, which relies on a central gateway being responsible for managing and handling citizen authentications. In order to mitigate the associated scalability issues, the PEPS approach could be moved into the public cloud. However, a move of a trusted service into the public cloud brings up new obstacles, especially with respect to citizens' privacy. In this paper we propose an approach how this move could be successfully realized by still preserving citizens' privacy and keeping existing national eID infrastructures untouched. We present the approach in detail and evaluate its capability with respect to citizens' privacy protection as well as its practicability. We conclude, that the proposed approach is a viable way of realizing an efficient and scalable Pan-European citizen identification and authentication framework.

- Markulf Kohlweiss and Alfredo Rial, "Optimally Private Access Control," in *Workshop on Privacy in the Electronic Society 2013 (WPES 2013)* [76].
  Abstract. Access control based on anonymous credentials allows users to prove to a service provider in a privacy-friendly manner that they possess the credentials required to access a resource. To achieve optimal privacy, the information that service providers can learn from the access control protocol should in principle be just a single event, namely that a user is granted access. However, existing anonymous credential schemes reveal additional information to the service provider such as the identity of the credential issuer, the credential type, and constraints on the attributes of the credential that reveal more than the access decision itself. In addition, the efficiency of selective attribute disclosure is not optimal.

  Our contribution is both cryptographic and conceptual. First, we extend existing vector commitment schemes with efficient zero-knowledge protocols to prove correct generation

want to conduct online procedures with their citizen cards.

of a new commitment, to prove that a secret value is committed at a secret position, and to prove that a secret position was updated to a new secret value. Second, we employ these protocols along with structure preserving signatures and conceptual techniques from logic-based access control to design a private access control protocol with efficient selective attribute disclosure that achieves our optimality criteria.

- Jan Camenisch and Robert Enderlein and Victor Shoup, "Practical and Employable Protocols for UC-Secure Circuit Evaluation over Z_n," in *ESORICS 2013* [30].
  Abstract. We present a set of new, efficient, universally composable two-party protocols for evaluating reactive arithmetic circuits modulo n, where n is a safe RSA[5] modulus of unknown factorization. Our protocols are based on a homomorphic encryption scheme with message space Z_n, zero-knowledge proofs of existence, and a novel "mixed" trapdoor commitment scheme. Our protocols are proven secure against adaptive corruptions (assuming secure erasures) under standard assumptions in the CRS model (without random oracles). Our protocols appear to be the most efficient ones that satisfy these security requirements. In contrast to prior protocols, we provide facilities that allow for the use of our protocols as building blocks of higher-level protocols. An additional contribution of this paper is a universally composable construction of the variant of the Dodis- Yampolskiy oblivious pseudorandom function in a group of order n as originally proposed by Jarecki and Liu.

- Masayuki Abe and Jan Camenisch and Maria Dubovitskaya and Ryo Nishimaki, "Universally Composable Adaptive Oblivious Transfer (with Access Control) from Standard Assumptions," in *Digital Identity Management Workshop 2013* [7].
  Abstract. In this paper, we provide the first scheme that realises an attribute-based access control system for static resources that offers maximal privacy and is secure in the universally composability framework (UC). More precisely, we offer a protocol for adaptive oblivious transfer where the sender can enforce an attribute-based access control policy for each record and nevertheless does not learn which record a user retrieves nor which attributes a user has. As additional results we provide a new structure-preserving signature scheme and a new universally composable adaptive oblivious transfer protocol that is secure under two DDH-like assumptions and is the most efficient one secure under non "q-type" assumptions[6]. We believe the new signature scheme to be of independent interest as a building block that is compatible with Groth-Sahai non-interactive zero-knowledge proofs.

- Changyu Dong and Liqun Chen and Jan Camenisch and Giovanni Russello, "Fair Private Set Intersection with a Semi-trusted Arbiter," [53].
  Abstract. A private set intersection (PSI) protocol allows two parties to compute the intersection of their input sets privately. Most of the previous PSI protocols only output the result to one party and the other party gets nothing from running the protocols. However, a mutual PSI protocol in which both parties can get the output is highly desirable in many applications. A major obstacle in designing a mutual PSI protocol is how to ensure

[5] RSA is an algorithm for public-key cryptography that is based on the presumed difficulty of factoring large integers.

[6] In "q-type" assumptions, the length of the input to the problem grows with "q".

fairness. In this paper we present the first fair mutual PSI protocol which is efficient and secure. Fairness of the protocol is obtained in an optimistic fashion, i.e. by using an offline third party arbiter. In contrast to many optimistic protocols which require a fully trusted arbiter, in our protocol the arbiter is only required to be semi-trusted, in the sense that we consider it to be a potential threat to both parties' privacy but believe it will follow the protocol. The arbiter can resolve disputes without knowing any private information belongs to the two parties. This feature is appealing for a PSI protocol in which privacy may be of ultimate importance.

- Jan Camenisch and Maria Dubovitskaya and Anja Lehmann and Gregory Neven and Christian Paquin and Franz-Stefan Preiss, "Concepts and Languages for Privacy-Preserving Attribute-Based Authentication," [29].
  Abstract. Existing cryptographic realizations of privacy-friendly authentication mechanisms such as anonymous credentials, minimal disclosure tokens, self-blindable credentials, and group signatures vary largely in the features they offer and in how these features are realized. Some features such as revocation or de-anonymization even require the combination of several cryptographic protocols. These differences and the complexity of the cryptographic protocols hinder the deployment of these mechanisms for practical applications and also make it almost impossible to switch the underlying cryptographic algorithms once the application has been designed. In this paper, we aim to overcome this issue and simplify both the design and deployment of privacy-friendly authentication mechanisms. We define and unify the concepts and features of privacy-preserving attribute-based credentials (Privacy-ABCs) and provide a language framework in XML schema. Our language framework enables application developers to use Privacy-ABCs with all their features without having to consider the specifics of the underlying cryptographic algorithms—similar to as they do today for digital signatures, where they do not need to worry about the particulars of the RSA and DSA[7] algorithms either.

- Jan Camenisch and Anja Lehmann and Gregory Neven and Alfredo Rial, "Privacy-Enhancing Audit Methods for Attribute-Based Credentials," [33].
  Abstract. Attribute-based credentials allow users to send to a verifier presentation tokens that prove that they possess certain attributes. For auditing purposes, in some settings the verifier reveals those presentation tokens to an auditor. In existing attribute-based credentials schemes, the verifier has to reveal the entire presentation token to the auditor, thus disclosing all the user's attributes that were revealed to the verifier. We provide an audit mechanism that allows the verifier to disclose only a subset of those attributes. Furthermore, in our scheme the user decides which attributes the verifier can forward to the auditor. We present two instantiations of our scheme based on Damgård-Fujisaki commitments and on polynomial commitments.

- Thomas Groß, "Signatures and Efficient Proofs on Committed Graphs and NP-Statements", under submission. [61].
  Abstract. Digital signature schemes are a foundational cryptographic building block enabling integrity and non-repudiation. We propose a graph signature scheme and corresponding proofs that allow a prover (1) to obtain a signature on a committed graph and

[7] The Digital Signature Algorithm (DSA) is a Federal Information Processing Standard for digital signatures.

(2) to subsequently prove to a verifier knowledge of such a graph signature. The graph signature scheme and proofs are a building block for certification systems that need to prove graph properties in zero-knowledge, as encountered in cloud security assurance or provenance. We extend the Camenisch-Lysyanskaya (CL) signature scheme to graphs and enable efficient zero-knowledge proofs of knowledge on graph signatures, notably supporting complex statements on graph elements. Our method is based on honest-verifier proofs and the strong RSA assumption. We show how the system enables cloud topology proofs and relates to Direct Anonymous Attestation. In addition, we explore the capabilities of graph signatures by establishing a proof system on graph 3-colorability (G3C). As G3C is NP-complete, we conclude that there exist Camenisch-Lysyanskaya proof systems for statements of NP languages.

10.4 Research on Privacy-Friendly Revocation Mechanisms (Task 24.4)

We do not have publications for this task so far.

10.5 Methods for Usable Privacy (Task 24.5)

We do not have publications for this task so far.

List of References

[1] GÉANT network. 2013. URL: http://www.geant.net/pages/home.aspx.

[2] SEMIRAMIS project. 2013. URL: www.semiramis-cip.eu.

[3] SWIFT project. 2013. URL: http://www.ist-swift.org/.

[4] Martín Abadi. Private authentication. In *PET*, LNCS 2482, pages 27–40. Springer, 2003.

[5] Martín Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In *POPL*, pages 104–115. ACM, 2001. doi:10.1145/360204.360213.

[6] Martín Abadi and Roger Needham. Prudent engineering practice for cryptographic protocols. *IEEE Trans. Softw. Eng.*, 22(1):6–15, January 1996. URL: http://dx.doi.org/10.1109/32.481513, doi:10.1109/32.481513.

[7] Masayuki Abe, Jan Camenisch, Maria Dubovitskaya, and Ryo Nishimaki. Universally Composable Adaptive Oblivious Transfer (with Access Control) from Standard Assumptions. In *Digital Identity Management Workshop 2013*. ACM, 2013.

[8] A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *Security Privacy, IEEE*, 3(1):26–33, 2005. doi:10.1109/MSP.2005.22.

[9] Anne Adams and Martina Angela Sasse. Users are not the enemy. *Commun. ACM*, 42(12):40–46, December 1999. URL: http://doi.acm.org/10.1145/322796.322806, doi:10.1145/322796.322806.

[10] Jae Hyun Ahn, Dan Boneh, Jan Camenisch, Susan Hohenberger, Abhi Shelat, and Brent Waters. Computing on authenticated data. *IACR Cryptology ePrint Archive*, 2011:96, 2011. URL: http://dblp.uni-trier.de/db/journals/iacr/iacr2011.html#AhnBCHSW11.

[11] Jae Hyun Ahn, Dan Boneh, Jan Camenisch, Susan Hohenberger, Abhi Shelat, and Brent Waters. Computing on Authenticated Data. In *TCC*, volume 7194 of *LNCS*, pages 1–20. Springer, 2012.

[12] Myrto Arapinis, Tom Chothia, Eike Ritter, and Mark Ryan. Analysing unlinkability and anonymity using the applied pi calculus. In *Proceedings of the 2010 23rd IEEE Computer Security Foundations Symposium*, CSF '10, pages 107–121, Washington, DC, USA, 2010. IEEE Computer Society. URL: http://dx.doi.org/10.1109/CSF.2010.15, doi:10.1109/CSF.2010.15.

[13] Myrto Arapinis and Marie Duflot. Bounding messages for free in security protocols. In *FSTTCS*, pages 376–387, 2007.

[14] Giuseppe Ateniese, Daniel H. Chou, Breno de Medeiros, and Gene Tsudik. Sanitizable signatures. In Sabrina De Capitani di Vimercati, Paul F. Syverson, and Dieter Gollmann, editors, *ESORICS*, volume 3679 of *Lecture Notes in Computer Science*, pages 159–177. Springer, 2005. URL: http://dblp.uni-trier.de/db/conf/esorics/esorics2005.html#AtenieseCMT05.

[15] Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Trans. Inf. Syst. Secur.*, 9(1):1–30, 2006.

[16] Ricardo Azevedo, Tiago Batista, Antonio F. Skarmeta Gómez, Gregorio Martín Pérez, Juan Manuel Marín Pérez, Jorge Bernal Bernabé, Manuel Gil Pérez, Aljosa Pasic, Charles Bastos Rodríguez, and Rubén Torres Diéguez. Privacy-preserving Identity Management in SEMIRAMIS. 2013. URL: `http://www.geant.net/pages/home.aspx`.

[17] Michael Backes and David A. Basin, editors. *Proceedings of the 2003 ACM workshop on Formal methods in security engineering, FMSE 2003, Washington, DC, USA, October 30, 2003*. ACM, 2003.

[18] Michael Backes, Sebastian Meiser, and Dominique Schröder. Delegatable Functional Signatures. *IACR Cryptology ePrint Archive*, 2013:408, 2013.

[19] David A. Basin, Sebastian Mödersheim, and Luca Viganò. Ofmc: A symbolic model checker for security protocols. *Int. J. Inf. Sec.*, 4(3):181–208, 2005.

[20] Mihir Bellare and Georg Fuchsbauer. Policy-Based Signatures. *IACR Cryptology ePrint Archive*, 2013:413, 2013.

[21] Abhilasha Bhargav-Spantzel, Jan Camenisch, Thomas Gross, and Dieter Sommer. User centricity: A taxonomy and open issues. *J. Comput. Secur.*, 15(5):493–527, October 2007. URL: `http://dl.acm.org/citation.cfm?id=1370624.1370625`.

[22] Bruno Blanchet. An efficient cryptographic protocol verifier based on prolog rules. In *IN 14TH IEEE COMPUTER SECURITY FOUNDATIONS WORKSHOP (CSFW-14*, pages 82–96. IEEE Computer Society Press, 2001.

[23] Bruno Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Trans. Dependable Sec. Comput.*, 5(4):193–207, 2008.

[24] Bruno Blanchet and Andreas Podelski. Verification of cryptographic protocols: tagging enforces termination. *Theor. Comput. Sci.*, 333(1-2):67–90, 2005.

[25] Dan Boneh and David Mandell Freeman. Homomorphic Signatures for Polynomial Functions. In *EUROCRYPT*, pages 149–168, 2011.

[26] Stefan Brands, Liesje Demuynck, and Bart De Decker. A practical system for globally revoking the unlinkable pseudonyms of unknown users. In *Proceedings of the 12th Australasian conference on Information security and privacy*, ACISP'07, pages 400–415, Berlin, Heidelberg, 2007. Springer-Verlag. URL: `http://dl.acm.org/citation.cfm?id=1770231.1770268`.

[27] Christina Brzuska, Marc Fischlin, Tobias Freudenreich, Anja Lehmann, Marcus Page, Jakob Schelbert, Dominique Schröder, and Florian Volk. Security of sanitizable signatures revisited. In *PKC '09*, LNCS, pages 317–336. Springer, 2009.

[28] Christina Brzuska, Marc Fischlin, Anja Lehmann, and Dominique Schröder. Unlinkability of sanitizable signatures. In Phong Q. Nguyen and David Pointcheval, editors, *Public Key Cryptography*, volume 6056 of *Lecture Notes in Computer Science*, pages 444–461. Springer, 2010. URL: http://dblp.uni-trier.de/db/conf/pkc/pkc2010.html#BrzuskaFLS10.

[29] Jan Camenisch, Maria Dubovitskaya, Anja Lehmann, Gregory Neven, Christian Paquin, and Franz-Stefan Preiss. Concepts and languages for privacy-preserving attribute-based authentication. In Simone Fischer-Hübner, Elisabeth de Leeuw, and Chris Mitchell, editors, *IDMAN*, volume 396 of *IFIP Advances in Information and Communication Technology*, pages 34–52. Springer, 2013.

[30] Jan Camenisch, RobertR. Enderlein, and Victor Shoup. Practical and employable protocols for uc-secure circuit evaluation over z_n. In Jason Crampton, Sushil Jajodia, and Keith Mayes, editors, *Computer Security – ESORICS 2013*, volume 8134 of *Lecture Notes in Computer Science*, pages 19–37. Springer Berlin Heidelberg, 2013. URL: http://dx.doi.org/10.1007/978-3-642-40203-6_2, doi:10.1007/978-3-642-40203-6_2.

[31] Jan Camenisch, Markulf Kohlweiss, and Claudio Soriente. An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In *Proceedings of the 12th International Conference on Practice and Theory in Public Key Cryptography: PKC '09*, Irvine, pages 481–500, Berlin, Heidelberg, 2009. Springer-Verlag. URL: http://dx.doi.org/10.1007/978-3-642-00468-1_27, doi:10.1007/978-3-642-00468-1_27.

[32] Jan Camenisch, Markulf Kohlweiss, and Claudio Soriente. Solving revocation with efficient update of anonymous credentials. In *Proceedings of the 7th international conference on Security and cryptography for networks*, SCN'10, pages 454–471, Berlin, Heidelberg, 2010. Springer-Verlag. URL: http://dl.acm.org/citation.cfm?id=1885535.1885576.

[33] Jan Camenisch, Anja Lehmann, Gregory Neven, and Alfredo Rial. Privacy-Enhancing Audit Methods for Attribute-Based Credentials. 2013.

[34] Jan Camenisch and Anna Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In *Proceedings of the 22nd Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '02, pages 61–76, London, UK, UK, 2002. Springer-Verlag. URL: http://dl.acm.org/citation.cfm?id=646767.704437.

[35] Jan Camenisch, Sebastian Mödersheim, Gregory Neven, Franz-Stefan Preiss, and Dieter Sommer. A card requirements language enabling privacy-preserving access control. In *Proceedings of the 15th ACM symposium on Access control models and technologies*, SACMAT '10, pages 119–128, New York, NY, USA, 2010. ACM. URL: http://doi.acm.org/10.1145/1809842.1809863, doi:10.1145/1809842.1809863.

[36] Kim Cameron and Michael B Jones. Design rationale behind the identity metasystem architecture. In *ISSE/SECURE 2007 Securing Electronic Business Processes*, pages 117–129. Springer, 2007.

[37] Sébastien Canard and Roch Lescuyer. Protecting privacy by sanitizing personal data: a new approach to anonymous credentials. In *ASIA CCS '13*, pages 381–392. ACM, 2013.

[38] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. Report 2000/067, Cryptology EPrint Archive, October 2001. Extended Abstract appeared in proceedings of the 42nd Symposium on Foundations of Computer Science (FOCS), 2001.

[39] Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Sarah Meiklejohn. Malleable Signatures: Complex Unary Transformations and Delegatable Anonymous Credentials. Cryptology ePrint Archive, Report 2013/179, 2013. `http://eprint.iacr.org/`.

[40] Vincent Cheval, Hubert Comon-Lundh, and Stéphanie Delaune. Automating security analysis: Symbolic equivalence of constraint systems. In Jürgen Giesl and Reiner Hähnle, editors, *IJCAR*, volume 6173 of *Lecture Notes in Computer Science*, pages 412–426. Springer, 2010. URL: `http://dblp.uni-trier.de/db/conf/cade/ijcar2010.html#ChevalCD10`.

[41] Stefan Ciobâca and Véronique Cortier. Protocol composition for arbitrary primitives. In *CSF*, pages 322–336. IEEE Computer Society, 2010.

[42] Stefan Ciobaca and Véronique Cortier. Protocol composition for arbitrary primitives. In *CSF*, pages 322–336. IEEE Computer Society, 2010. URL: `http://dblp.uni-trier.de/db/conf/csfw/csf2010.html#CiobacaC10`.

[43] Véronique Cortier, Jérémie Delaitre, and Stéphanie Delaune. Safely composing security protocols. In V. Arvind and Sanjiva Prasad, editors, *Proceedings of the 27th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'07)*, volume 4855 of *Lecture Notes in Computer Science*, pages 352–363, New Delhi, India, December 2007. Springer. `doi:10.1007/978-3-540-77050-3_29`.

[44] Véronique Cortier and Stéphanie Delaune. Safely composing security protocols. *Formal Methods in System Design*, 34(1):1–36, 2009.

[45] Véronique Cortier, Michaël Rusinowitch, and Eugen Zalinescu. Relating two standard notions of secrecy. *Logical Methods in Computer Science*, 3(3), 2007.

[46] Kenneth James Williams Craik. *The nature of explanation*. CUP Archive, 1967.

[47] Anupam Datta, Ante Derek, John C. Mitchell, and Dusko Pavlovic. Secure protocol composition. In Backes and Basin [17], pages 11–23.

[48] Anupam Datta, Ante Derek, John C. Mitchell, and Dusko Pavlovic. Secure protocol composition. In Backes and Basin [17], pages 11–23.

[49] Björn Deiseroth, Victoria Fehr, Marc Fischlin, Manuel Maasz, Nils Fabian Reimers, and Richard Stein. Computing on Authenticated Data for Adjustable Predicates. In *ACNS*, LNCS. Springer, 2013.

[50] Stéphanie Delaune, Steve Kremer, and Mark Ryan. Verifying privacy-type properties of electronic voting protocols. *J. Comput. Secur.*, 17(4):435–487, December 2009. URL: `http://dl.acm.org/citation.cfm?id=1576303.1576305`.

[51] Stéphanie Delaune, Mark D. Ryan, and Ben Smyth. Automatic verification of privacy properties in the applied pi-calculus. In *IFIPTM*, pages 263–278. Springer, 2008.

[52] Tim Dierks and Christopher Allen. RFC 2246: The TLS protocol, January 1999. Status: Standards Track. URL: `ftp://ftp.rfc-editor.org/in-notes/rfc2246.txt`.

[53] Changyu Dong, Liqun Chen, Jan Camenisch, and Giovanni Russello. Fair private set intersection with a semi-trusted arbiter. *IACR Cryptology ePrint Archive*, 2012:252, 2012. informal publication. URL: `http://dblp.uni-trier.de/db/journals/iacr/iacr2012.html#DongCCR12`.

[54] Cynthia Dwork. Differential privacy: a survey of results. In *Proceedings of the 5th international conference on Theory and applications of models of computation*, TAMC'08, pages 1–19, Berlin, Heidelberg, 2008. Springer-Verlag. URL: `http://dl.acm.org/citation.cfm?id=1791834.1791836`.

[55] Cynthia Dwork. Differential Privacy: A Survey of Results. In *TAMC*, LNCS 4978, pages 1–19. Springer, 2008.

[56] Simone Fischer-Hübnner, Christina Köffel, John-Sören Pettersson, Peter Wolkerstorfer, Cornelia Graf, Leif-Erik Holtz, Ulrich König, Hans Hedbom, and Benjamin Kellermann. HCI pattern collection version 2. 2011.

[57] Cornelia Graf, Christina Hochleitner, Peter Wolkerstorfer, Julio Angulo, Simone Fischer-Hübner, and Erik Wästlund. Towards Usable Privacy Enhancing Technologies: Lessons Learned from the PrimeLife Project. 2011.

[58] Cornelia Graf, Peter Wolkerstorfer, Arjan Geven, and Manfred Tscheligi. A pattern collection for privacy enhancing technology. In *PATTERNS 2010, The Second International Conferences on Pervasive Patterns and Applications*, pages 72–77, 2010.

[59] Cornelia Graf, Peter Wolkerstorfer, Christina Hochleitner, Erik Wästlund, and Manfred Tscheligi. Hci for primelife prototypes. In Jan Camenisch, Simone Fischer-Hübner, and Kai Rannenberg, editors, *Privacy and Identity Management for Life*, pages 221–232. Springer Berlin Heidelberg, 2011. URL: `http://dx.doi.org/10.1007/978-3-642-20317-6_11`, `doi:10.1007/978-3-642-20317-6_11`.

[60] Matthew Green and Giuseppe Ateniese. Identity-Based Proxy Re-encryption. In *ACNS*, volume 4521 of *LNCS*, pages 288–306. Springer, 2007.

[61] Thomas Groß. Signatures and efficient proofs on committed graphs and np-statements. Under submission.

[62] Thomas Groß and Sebastian Mödersheim. Vertical protocol composition. In *24th IEEE Computer Security Foundations Symposium (CSF) 2011*, pages 235–250. IEEE, 2011.

[63] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel P. Smart, editor, *EUROCRYPT*, volume 4965 of *Lecture Notes in Computer Science*, pages 415–432. Springer, 2008. URL: `http://dblp.uni-trier.de/db/conf/eurocrypt/eurocrypt2008.html#GrothS08`.

[64] Thomas Groß and Sebastian Mödersheim. Vertical protocol composition. In *Proceedings of the 24th IEEE Computer Security Foundations Symposium, CSF 2011, Cernay-la-Ville, France, 27-29 June, 2011*, pages 235–250. IEEE Computer Society, 2011. doi:http://doi.ieeecomputersociety.org/10.1109/CSF.2011.23.

[65] Joshua D. Guttman. Cryptographic protocol composition via the authentication tests. In Luca de Alfaro, editor, *FOSSACS*, volume 5504 of *Lecture Notes in Computer Science*, pages 303–317. Springer, 2009. URL: http://dblp.uni-trier.de/db/conf/fossacs/fossacs2009.html#Guttman09.

[66] Joshua D. Guttman and F. Javier Thayer. Protocol independence through disjoint encryption. In *CSFW*, pages 24–34, 2000.

[67] Christian Hanser and Daniel Slamanig. Blank Digital Signatures. In *8th ACM SIGSAC Symposium on Information, Computer and Communications Security (AsiaCCS)*, pages 95–106. ACM, 2013.

[68] Christian Hanser and Daniel Slamanig. Warrant-Hiding Delegation-by-Certificate Proxy Signature Schemes. In *14th International Conference on Cryptology in India (INDOCRYPT)*, volume 8250 of *LNCS*. Springer, 2013. to appear.

[69] James Heather, Gavin Lowe, and Steve Schneider. How to prevent type flaw attacks on security protocols. In *CSFW*, pages 255–268, 2000.

[70] Timothy Hinrichs and Michael Genesereth. Herbrand logic. Technical Report LG-2006-02, Stanford University, CA, USA, 2006. http://logic.stanford.edu/reports/LG-2006-02.pdf.

[71] Leif-Erik Holtz, Harald Zwingelberg, and Marit Hansen. Privacy policy icons. In *Privacy and Identity Management for Life*, pages 279–285. Springer, 2011.

[72] IBM Research – Zurich. Specification of the identity mixer cryptographic library. version 2.3.4. Technical report, IBM Research, 2012.

[73] Robert Johnson, David Molnar, Dawn Xiaodong Song, and David Wagner. Homomorphic Signature Schemes. In *CT-RSA*, volume 2271 of *LNCS*, pages 244–262. Springer, 2002.

[74] Philip N Johnson-Laird and Ruth MJ Byrne. Conditionals: a theory of meaning, pragmatics, and inference. *Psychological review*, 109(4):646, 2002.

[75] Philip Nicholas Johnson-Laird. *Mental models: Towards a cognitive science of language, inference, and consciousness*, volume 6. Harvard University Press, 1983.

[76] Markulf Kohlweiss and Alfredo Rial. Optimally Private Access Control. In *12th Workshop on Privacy in the Electronic Society*. ACM, 2013.

[77] Ninghui Li and Tiancheng Li. t-closeness: Privacy beyond k-anonymity and l-diversity. In *In Proc. of IEEE 23rd Int'l Conf. on Data Engineering (ICDE'07*, 2007.

[78] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubramaniam. L-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1), March 2007. URL: http://doi.acm.org/10.1145/1217299.1217302, doi:10.1145/1217299.1217302.

[79] Gregorio Martínez. End-to-end encryption and Attribute translation features. Demo of Apply for Courses use case. 2013. URL: http://www.geant.net/pages/home.aspx.

[80] Sebastian Mödersheim. Deciding security for a fragment of aslan. In *ESORICS*, pages 127–144, 2012.

[81] Sebastian Mödersheim and Luca Viganò. Secure pseudonymous channels. In Michael Backes and Peng Ning, editors, *ESORICS*, volume 5789 of *Lecture Notes in Computer Science*, pages 337–354. Springer, 2009.

[82] Sebastian Alexander Mödersheim. Abstraction by set-membership: verifying security protocols and web services with databases. In *Proceedings of the 17th ACM conference on Computer and communications security*, CCS '10, pages 351–360, New York, NY, USA, 2010. ACM. URL: http://doi.acm.org/10.1145/1866307.1866348, doi:10.1145/1866307.1866348.

[83] Toru Nakanishi, Hiroki Fujii, Yuta Hira, and Nobuo Funabiki. Revocable group signature schemes with constant costs for signing and verifying. In *Proceedings of the 12th International Conference on Practice and Theory in Public Key Cryptography: PKC '09*, Irvine, pages 463–480, Berlin, Heidelberg, 2009. Springer-Verlag. URL: http://dx.doi.org/10.1007/978-3-642-00468-1_26, doi:10.1007/978-3-642-00468-1_26.

[84] Jakob Nielsen. *Usability Engineering*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1993.

[85] Donald A. Norman. *The Design of Everyday Things*. Basic Books, New York, reprint paperback edition, 2002.

[86] Birgit Pfitzmann and Michael Waidner. A model for asynchronous reactive systems and its application to secure message transmission. In *Security and Privacy, 2001. S&P 2001. Proceedings. 2001 IEEE Symposium on*, pages 184–200. IEEE, 2001.

[87] Bart Priem, Eleni Kosta, Aleksandra Kuczerawy, Jos Dumortier, and Ronald Leenes. User-centric privacy-enhancing identity management. In Jan Camenisch, Ronald Leenes, and Dieter Sommer, editors, *Digital Privacy*, volume 6545 of *Lecture Notes in Computer Science*, pages 91–106. Springer Berlin Heidelberg, 2011. URL: http://dx.doi.org/10.1007/978-3-642-19050-6_6, doi:10.1007/978-3-642-19050-6_6.

[88] Jerome H. Saltzer and Michael D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, 1975. URL: http://dblp.uni-trier.de/db/journals/pieee/pieee63.html#SaltzerS75.

[89] Ron Steinfeld, Laurence Bull, and Yuliang Zheng. Content Extraction Signatures. In *ICISC*, volume 2288 of *LNCS*, pages 163–205. Springer, 2002.

[90] Latanya Sweeney. k-anonymity: a model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570, October 2002. URL: `http://dx.doi.org/10.1142/S0218488502001648`, `doi:10.1142/S0218488502001648`.

[91] Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):557–570, 2002.

[92] S. Trabelsi, J. Sendor, and S. Reinicke. Ppl: Primelife privacy policy engine. In *Policies for Distributed Systems and Networks (POLICY), 2011 IEEE International Symposium on*, pages 184–185, 2011. `doi:10.1109/POLICY.2011.24`.

[93] Erik Wästlund and Simone Fischer-Hübner. The users' mental models' effect on their comprehension of anonymous credentials. In *Privacy and Identity Management for Life*, pages 233–244. Springer, 2011.

[94] C. Weidenbach, R. A. Schmidt, T. Hillenbrand, R. Rusev, and D. Topic. System description: SPASS version 3.0. In F. Pfenning, editor, *Automated Deduction—CADE-21*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 514–520. Springer, 2007. `doi:http://dx.doi.org/10.1007/978-3-540-73595-3_38`.

[95] Bernd Zwattendorfer and Daniel Slamanig. On Privacy-Preserving Ways to Porting the Austrian eID System to the Public Cloud. In Sujeet Shenoi Lech J. Janczewski, Henry B. Wolfe, editor, *28th IFIP TC-11 International Information Security and Privacy Conference (SEC 2013)*, volume 405 of *IFIP AICT*, pages 300–314. Springer, 2013.

[96] Bernd Zwattendorfer and Daniel Slamanig. Privacy-Preserving Realization of the STORK Framework in the Public Cloud. In *10th International Conference on Security and Cryptography (SECRYPT 2013)*, pages 419–426, 2013.