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

D24.4

Document Identification	
Date	October 27, 2015
Status	Final
Version	1.0

Related SP/WP	SP2/WP24	**Document Reference**	D24.4
Related Deliverable(s)	D12.3, D12.4, D22.1, D22.2, D22.3, D23.1, D24.1, D24.2, D24.3, 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) Paolo Modesti (UNEW) Daniel Slamanig (TUG) Sebastian Mödersheim (DTU) Omar Almousa (DTU) Jaap-Henk Hoepman (RU)	**Reviewers**	Lothar Fritsch (NRS) Tobias Wich (ECS)

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 the FutureID project is to build a comprehensive, flexible, privacy-aware and ubiquitously usable identity management infrastructure for Europe. That is, the main goal of the project is to provide an architecture that allows the different eID solutions already deployed to be used in a unified and interoperable manner.

The aim of work package 24 is to address various shortcomings 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 third year of the FutureID project and provides a summary of the research results.

Task 24.1: Extending languages and tools for compositional reasoning. The purpose of this task is to provide languages and tools that allow the analysis of complex systems that are composed of multiple components. We have achieved a significant step forward in the area of compositional reasoning with two kinds of *relative soundness results*. The first kind are typing results showing that any security protocol that fulfils a number of sufficient conditions has an attack if it has a well-typed attack. The second kind considers the parallel composition of protocols, showing that when running two protocols in parallel allows for an attack, then at least one of the protocols has an attack in isolation. In this deliverable, we present 2 publications related to this task. In WP24, we have presented 5 publications.

Task 24.2: Establishing methods and languages for privacy goals. This task's goal is to establish 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. For the first sub-task, we focus on applying the concept of α-β-privacy that we have developed previously to the FutureID architecture. For the second sub-task, we define and unify the concepts and features of privacy-preserving attribute-based credentials (Privacy-ABCs), provide a language framework in XML schema, and give a formal semantics to describe the effects of the transactions in a privacy-friendly authentication system using Privacy-ABCs. Additionally, we present a Prolog implementation for credential-policy matching. In this deliverable, we present 1 publication and 2 technical reports related to this task. In WP24, we have presented 2 publications and 2 technical reports.

Task 24.3: Privacy-friendly audit and data-handling mechanisms. Task 24.3 conducts research on data-handling mechanisms and on audits. For audits, we report on experiences during implementing blank digital signatures as well as optimizations that helped to improve their performance. Additionally, we propose a novel graph signature scheme, which makes it possible that an issuer certifies a committed graph, such that a prover

can subsequently prove properties of the graph in zero-knowledge proofs of knowledge. Data-handling mechanisms determine how user data is managed by the service provider. We focus our research on authentication mechanisms based on passwords and signatures and on privacy-preserving protocols that minimize the data that users have to disclose to service providers. We also work on existing eID solutions. In addition, we also conduct research on computations on signed data and on data anonymization and data sharing between databases. In this deliverable, we present 12 publications and 1 technical report related to this task. In WP24, we have presented 27 publications and 1 technical report.

Task 24.4: Development of privacy-friendly revocation mechanisms. We address the design of several privacy-friendly revocation mechanisms. First, we propose a privacy preserving revocation mechanism for privacy-enhancing attribute-based credentials that allows you to efficiently handle multiple revocation lists. Second, we study a primitive that is widely used for revocation purposes, i.e., cryptographic accumulators. Third, we show how using epochs can help to make revocation practical while still retaining reasonable strong privacy guarantees. Our contribution is a new revocation scheme that has very low computational cost for users and verifiers alike, that is efficient even in the smart card setting, and therefore can be used in practice. Finally, we explain the concept of revocable privacy. In this deliverable, we present 3 publications and 1 technical report related to this task. In WP24, so far we have presented 4 publications and 1 technical report.

Task 24.5: Development of methods for usable privacy. We have designed a two-factor user-authentication scheme for usable server-based eID and e-signature solutions. Current server-based eID and e-signature solutions typically rely on one-time passwords delivered to the user via short message service (SMS). This raises several issues in practice, as the use of SMS technology can be cost-effective insecure. To address these issues, we propose an alternative two-factor user-authentication scheme following a challenge-response approach. The feasibility and applicability of the proposed user-authentication scheme is evaluated by means of two concrete implementations. This way, we show that the proposed authentication scheme and its implementations improve both the cost effectiveness and the security of server-based eID and e-signature solutions. Additionally, on a different line of work, we study how users choose passwords under consideration of different human dimensions, and, more specifically, when they are cognitively depleted. In this deliverable, we present 1 publication related to this task. In WP24, so far we have presented 5 publications and 1 poster.

2 Document information

2.1 Contributors

Name	**Partner**
Jan Camenisch	IBM
Alfredo Rial	IBM
Thomas Groß	UNEW
Paolo Modesti	UNEW
Daniel Slamanig	TUG
Sebastian Mödersheim	DTU
Omar Almousa	DTU
Jaap-Henk Hoepman	RU

2.2 History

0.01	2015-08-31	Alfredo Rial	1st Draft
0.02	2015-09-04	Alfredo Rial	IBM contribution tasks 24.2, 24.3 and 24.4
0.03	2015-09-25	Paolo Modesti	UNEW contribution task 24.1
0.04	2015-10-08	Jaap-Henk Hoepman	RU contribution task 24.4
0.05	1015-10-16	Thomas Groß	UNEW contribution tasks 24.3 and 24.5
0.06	2015-10-18	Omar Almousa	DTU contribution task 24.2
0.07	2015-10-18	Sebastian Mödersheim	DTU contribution task 24.2
0.08	2015-10-19	Alfredo Rial	Executive summary and conclusion
0.09	2015-10-19	Alfredo Rial	Front page and bibliography
1.0	2015-10-27	Alfredo Rial	Addressed reviewers' comments

2.3 Table of Contents

2.4 List of Figures

2.5 List of Tables

2.6 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.

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.

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

3.1 Document Structure

In Section 3, we first recall the Description of Work for WP24 in Section 3.2 and then we summarize the research conducted for WP24 during the three years of the project in Section 3.3. In Section 4, Section 5, Section 6, Section 7, and Section 8, we describe the research conducted during the third year of the project for Task 24.1, Task 24.2, Task 24.3, Task 24.4 and Task 24.5 respectively. We conclude in Section 9. In Section 10, we include the abstracts of all the publications described in this document, while in Section 11 we include the references of all the publications described in D24.1, D24.2, and D24.4. (The publications in D24.3 are also described in D24.4.) Therefore, Section 10 contains the abstracts of the publications presented in WP24 during the third year of the project, while Section 11 contains the references of the publications presented in WP24 during the whole project.

3.2 Description of Work for WP24

In the following subsections, we recall the research goals described in the Description of Work of FutureID for each of the tasks.

3.2.1 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 On-the-Fly Model Checker (OFMC) [44], AIF/ProVerif/SPASS tool-chain [268], [60], [341], which are complex systems.

They are composed of 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 [158], [201], [155], [148], [195] 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 [8] 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.

3.2.2 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 [145], [162], [29]. 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 [326] (which, in a database, suggests the suppression and generalization (obfuscation) of quasi-identifiers to make an individual's data entry indistinguishable from others) and its derivatives [251], [242], and differential privacy [169].

Task 24.2.2: Formal Semantics for Claims Languages. This task develops a formal semantics for the claim-language following the example of the card requirements language CARL [120]. 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 were made by the same user or not, unless users deliberately link their actions.

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

Users reveal data to service providers for authentication and authorization 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 [329]. 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 [32], [83]. We will investigate how current privacy-enhanced credentials and eID solutions can be extended to allow for such legitimate post-processing as well.

3.2.4 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 [78], [115], [104], [105], [272] 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.

3.2.5 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" [312] 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 [58] and realized in [125] 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 [280] and therefore satisfy qualitative properties including the following ones:

- Learnability, efficiency, memorability, low error rates and high satisfaction [280].
- Control [282], while considering limiting factors such as information, time, and psychological deviations [21].

These limitations dictate that citizens may not be able to fully understand the different threats and risks [22] 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.

3.3 Summary of the Research Conducted in WP24

We summarize the number of publications and technical reports in work package 24 in Table 1. The table includes publications and reports published in D24.1, D24.2 and in the present document. In the following, we summarize the research topics addressed by those publications and reports. The reference list for those papers can be found in Section 11.

Table 1: Number of publications in work package 24.

	Pub.	Rep.
Task 24.1	5	0
Task 24.2	2	2
Task 24.3	27	1
Task 24.4	4	1
Task 24.5	5	1

3.3.1 Summary of the Research Conducted in Task 24.1

In Task 24.1, the goal is to provide languages and tools that allow the analysis of complex systems that are composed of multiple components. We have carried out both theoretical and practical research for this task. For the theoretical part, we have achieved a significant step forward in the area of compositional reasoning with two kinds of *relative soundness results* [266, 267, 26]. For the practical part, we have created a tool for automatic generation of implementations of security protocols specified in a simple and abstract model that can be formally verified [270, 269].

3.3.2 Summary of the Research Conducted in Task 24.2

In Task 24.2, the goal is to establish methods and languages for the analysis of privacy goals with formal methods tools. Task 42.2 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. For the first sub-task, we have created the concept of α-β-privacy as privacy analysis method [265] and we have applied α-β-privacy to formally verify the security and privacy properties of the FutureID architecture [263]. For the second sub-task, we define and unify the concepts and features of privacy-preserving attribute-based credentials (Privacy-ABCs), provide a language framework in XML schema, and give a formal semantics to describe the effects of the transactions in a privacy-friendly authentication system using Privacy-ABCs [90]. Additionally, although not directly related with the goals of this task, we also describe a Prolog implementation for credential-policy matching [119].

3.3.3 Summary of the Research Conducted in Task 24.3

Task 24.3 conducts research on audits and on data-handling mechanisms. For audits, we have the following results. (Each of the items in the lists represents a separate publication.)

Attribute-Based Credentials. We have presented a scheme for privacy-preserving auditing of attribute-based credentials [113].

Proxy signatures. We have the following results.

- We propose warrant-hiding proxy signatures (WHPS) [207]. 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 not learning anything about the remaining message space.

- We propose blank digital signatures (BDS) [206]. 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.
- We provide black-box constructions of WHPS and BDS from non-interactive anonymous credentials [164].
- We report on experiences during implementing blank digital signatures as well as optimizations that helped to improve their performance [163].

Certification of committed graphs. We present the following results.

- We investigate efficient cryptographic primitives to certify the structure of a topology and to subsequently prove properties of the topology in zero-knowledge proofs of knowledge [196]. The primitive is highly applicable to identity federation topologies and uses FutureID-related components to facilitate the topology audit.
- We propose a novel graph signature scheme, which makes it possible that an issuer certifies a committed graph, such that a prover can subsequently prove properties of the graph in zero-knowledge proofs of knowledge [197].

For data-handling mechanisms, we have carried out research on different topics.

Password-Based Authentication. We present the following results.

- We describe a threshold password authenticated secret sharing protocol [97].
- We present two simple and extremely efficient proactively secure *distributed password verification* protocols [112].

Signature-Based Authentication. We present the following results.

- We give formal security definitions for a full-fledged privacy attribute-based credentials system. We provide a generic construction from lower-level building blocks that satisfies our definitions and we present secure instantiations of the building blocks [107].
- We propose a new kind of signature schemes, *unlinkable redactable block-signature* (URS) schemes, with which one can redact a signature and reveal only its relevant parts each time it is used. We construct an efficient URS scheme and we employ it to design the first universally composable anonymous credential system. It is also

arguably one of the first such schemes to support efficient attribute disclosure with cost independent of the number of attributes in the issued credential without having to rely on random oracles [92].

- We present an attribute-based credentials system that is optimally private, i.e., verifiers only learn whether users should be granted access or not, but no information on the user's attributes, the credential issuer or the policy the user fulfills [233].
- We present an efficient construction of round-optimal blind signature schemes in the standard model [183].
- We propose a "hybrid" group signature scheme, where unforgeability holds under classical assumptions, while privacy is proved under lattice-based ones. This allows us to combine the flexible tools that are available in the classical framework with the strong privacy guarantees of lattice problems. Our group signature scheme has keys and signatures of size logarithmic in the number of group members [56].
- We provide a novel type of structure-preserving signatures defined on equivalence classes on group element vectors, a novel randomizable polynomial commitment scheme, which allows to open factors of the polynomial committed to, and a new construction (type) of multi-show attribute-based anonymous credentials (ABCs), which is instantiated from the first two contributions [208].

Privacy-Preserving Protocols. We present the following results.

- We contribute a practical, secure and privacy-preserving mechanism enabling a service provider to verify whether the mobile phone of a given user currently resides within a certain geographical reference area at a given time. Our mechanism consists of having the location of the mobile phone determined by the Mobile Network Operator and certified using anonymous credentials [123].
- We propose a non-restricted and a restricted oblivious transfer with access control scheme [304].
- We provide the first oblivious transfer with access control protocol that is provably secure in the universal composability framework [12].
- We revisit existing work on privacy-preserving billing. First, we generalize the security model to consider multiple meters and multiple users. Second, we propose a privacy-preserving billing protocol for our model that improves efficiency for policies described by splines [306].
- We present the first fair mutual private set intersection protocol [168].
- We present the first solutions that enable social network users to share their externally hosted resources with social network friends while retaining a maximum level of privacy with respect to the service provider and the social network [102].

Computations on signed data. We show how the service provider can perform computations on unencrypted signed data [23].

Data anonymization and sharing. We propose an (un)linkable pseudonym system to allow exchange of user data between databases [110].

eID Solutions. We present the following results.

- Within the FutureID project, as use-case, we investigated (parts of) the Austrian eID ecosystem. 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) [350].
- We have applied the same kind of research done for the Austrian eID ecosystem to the Secure Identity Across Borders Linked (STORK) project [351].
- We present design strategies for a privacy-friendly Austrian eID system in the public cloud [352].
- We propose a novel and practical identification and authentication model to be applied for eIDs, which keeps the advantages of user-centricity but allows for selective disclosure possibilities to better protect citizens' privacy compared to existing national eID solutions [321].

Secure two-party computation. 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 [98].

3.3.4 Summary of the Research Conducted in Task 24.4

In Task 24.4, we address the design of several privacy-friendly revocation mechanisms. We propose a pairing-based group signature scheme with controllable linkability [320]. We also explain the concept of revocable privacy [247]. The rest of our research is centered on revocation for attribute-based credentials. First, we propose a privacy preserving revocation mechanism for privacy-enhancing attribute-based credentials that allows you to efficiently handle multiple revocation lists [96]. Second, we study a primitive that is widely used for revocation purposes, i.e., cryptographic accumulators [165]. Finally, we show how using epochs can help to make revocation practical while still retaining reasonable strong privacy guarantees. Our contribution is a new revocation scheme that has very low computational cost for users and verifiers alike, that is efficient even in the smart card setting, and therefore can be used in practice [246].

3.3.5 Summary of the Research Conducted in Task 24.5

The goal of Task 24.5 is to establish 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" their personal information is disclosed, as well as their awareness and express consent on any disclosure.

In Task 24.5, we conduct foundational research in usable privacy. This is mainly centred around empirical and scientific investigations of user privacy cognition via two main approaches: cognitive effort [153] and mental models [151]. Our research so far shows that comparing cognitive effort spent in privacy decision-making with other standardised effortless and effortful tasks will provide deeper understanding of the un-usability of privacy problems and potentially also help to explain the dichotomy between privacy attitudes and behaviour [152]. In addition, our mental models research provides initial of possible user segmentation and cognitive ability [149]. However we believe the methodology for the elicitation of user mental models of privacy needs to be firmly grounded in empirical methods due to the mental models uncertainty principle (that is that mental models are not directly accessible and observable and would vary with elicitation approach) [150]. Finally, in this deliverable, we study how users choose passwords under consideration of different human dimensions, and, more specifically, when they are cognitively depleted.

On a different line of work, we have designed a two-factor user-authentication scheme for usable server-based eID and e-signature solutions [302]. Additionally, we review the findings and recommendations on the usability of privacy enhancing identity management systems from the PRIME and PrimeLife projects. We also provide a case-study of Secure Identity Accross Borders Linked (STORK) with its privacy assessments and considerations for FutureID.

4 Extending Languages and Tools for Compositional Reasoning

In a collaboration between DTU (for the theoretical part) and UNEW (for the practical part) we have achieved a significant step forward in the area of compositional reasoning, which is fully described in [26] (an extended version is available in [27]).

This paper mainly focusses on two kinds of *relative soundness results*. This means we can give a general principle to reduce complex protocol verification problems to simpler ones, where the reduction ensures that if the complex protocol has an attack, then also the simple one has.[1] Thus, verifying that the simple protocol has no attack implies also the complex one is secure; it is thus sufficient to verify the simple protocol which is often considerably easier for automated methods, and also corresponds to the reasoning of the developer. Note that these soundness results will generally have requirements on the protocols for the reduction to be sound; these requirements are often called *sufficient conditions* (since protocols that do not satisfy them are not necessarily flawed) and they are of a "syntactical" nature, i.e., that can be checked by analyzing the structure of the protocol without having to explore the full state space.

The first kind of relative soundness results that we consider are *typing results* [210, 61, 264, 30]: here we go from a complex *untyped* model—where the intruder may introduce ill-typed messages to provoke type-flaw attacks—to a simpler typed model that is obtained by forbidding all ill-typed messages. While in general this is not a sound restriction (as demonstrated by the common type-flaw attacks), one can prove this sound under some restrictions on the message format. Note that the typed model may seem unreasonable at first sight, since in the real-world agents have no way to tell the type of a random bitstring, let alone distinguish it from the result of a cryptographic operation; yet in the model, they "magically" accept only well-typed messages. The relative soundness of such a typed model means that if the protocol has an attack, then it also has a well-typed attack. This does not mean that in the untyped model (where the intruder is able to send ill-typed messages) he cannot perform any attack that would not work similarly in the typed model. Thus, if we are able to verify that a protocol is secure in the typed model, then it is secure also in an untyped model. Typically, the conditions sufficient to achieve such a result are that all composed message patterns of the protocol have a different (intended) type that can somehow be distinguished, e.g., by a tag. The restriction to a typed model in some cases yields a decidable verification problem, allows for the application of more tools and often significantly reduces verification time in practice [61, 31].

The second kind of relative soundness results we consider is for *parallel composition* of protocols, i.e., running two protocols over the same communication medium, and these protocols may use, e.g., the same long-term public keys. (In the case of disjoint cryptographic material, compositional reasoning is relatively straightforward.) The compositionality result means to show that if two protocols satisfy their security goals in isolation, then their parallel composition is secure, provided the protocols meet certain sufficient conditions. Thus, it suffices to verify the protocols in isolation. The sufficient conditions in this case are similar to the typing result: every composed message can be uniquely attributed to one of the two protocols, which again may be

[1] The other direction is not necessarily true: the simple protocol may have an attack even though the complex one is fine.

achieved, e.g., by tags.

Our work unifies and thereby simplifies existing results on typing and parallel composition: we recast them as an instance of the same basic principle and of the same proof technique. In a nutshell, this technique is to reduce the search for attacks to solving constraint reduction in a symbolic model. For protocols that satisfy the respective sufficient conditions, the constraint reduction will never make an ill-typed substitution, where for compositionality "ill-typed" means to unify protocol messages from two different protocols.

This also allows us to significantly generalize existing results to a larger set of protocols and security properties. For what concerns protocols, our soundness results do not require a particular fixed tagging scheme like most previous works, but use more liberal requirements that are satisfied by many existing real-world protocols like TLS. While many existing results are limited to simple secrecy goals, we prove our results for the entire geometric fragment suggested by Guttman [202]. We even augment this fragment with the ability to directly refer to the intruder knowledge in the antecedent of goals; while this does not increase expressiveness, it is very convenient in specifications. In fact, handling the geometric fragment also constitutes a slight generalization of existing constraint-reduction approaches.

Another advantage over existing compositionality results is that to some extend we can even have compositionality with insecure protocols, while most other compositionality results require *all* composed protocols to be secure in isolation, for the composition to be secure. What we require instead is only that the protocols satisfy the sufficient conditions like disjointness of message formats and that they do not leak long-term secrets. These are properties that are usually very easy to verify for the protocols. For instance we can easily see that a protocol never leaks private keys, when these keys are only used for decryption and signing, but never sent as part of a decipherable message. The sufficient conditions, as said, are of the form that the message formats are disjoint, so that messages of one protocol cannot be parsed accidentally as those of another. Given that, we have the result that a secure protocol P_1 can be composed with any other (possibly insecure) protocol P_2 that at least has disjoint messages and never leaks common long-term secrets, then our compositionality result ensures that the goals of P_1 are never endangered by P_2.

To make these results available in practice to protocol designers and engineers, we have slightly extended our Authentication Protocol Specification Language SPS (defined in deliverable D42.3 and [25]) with a means to specify the necessary details for composition, namely which of the long-term keys have to be secret, and which have to be public, respectively. This is because one of the sufficient conditions for parallel composition is that the protocols agree on this, i.e., what is considered public in one protocol cannot be considered secret in another. With this, we have integrated into the APS translator (that can generate both protocol implementations and formal models) the full checking procedure for the sufficient conditions of typing and parallel composition[2]. A developer can thus feed a set of APS specifications into the APS translator and see if they are safe for parallel composition. If not, the developer gets a hint from the checker where a sufficient condition is violated and how the protocol could be changed in order to become composable.

[2]The tool is available at `http://www.imm.dtu.dk/~samo/SPS.zip`.

Finally, we want to mention that the paper received the best student paper award at ESORICS 2015, attributed to main author Omar Almousa whose PhD study has been mostly financed by FutureID.

AnBx At UNEW we extended our previous work [270] on the AnBx compiler [269], a tool for automatic generation of Java implementations of security protocols specified in a simple and abstract model that can be formally verified. In our model-driven development approach, protocols are described in AnBx, an extension of the Alice & Bob notation; along with the synthesis of consistency checks, the tool analyses the security goals and produces annotations that allow the verification of the generated implementation with ProVerif.

AIF-ω At DTU we have just finished an extension of the AIF verification tool that plays a crucial role in the formal verification of FutureID. This system is built on top of the popular ProVerif verification tool that is based on an abstraction approach: one over-approximates what can happen in a protocol by a set of Horn clauses and can then check whether an attack predicate is derivable from the Horn clauses by resolution. The benefit of this technique is to completely avoid the state-explosion problem of comon approaches, and verify protocols for an unbounded number of sessions. The drawback is that it only works for "monotonic" protocols. An example would be key revocation: what has been possible with that key before the revocation is not possible after. This is completely at odds with the Horn clause approach, since something that is true cannot become false when learning new information. To overcome this problem without destroying the original idea of abstraction, AIF was devised by DTU. The idea is that the modeler can declare a fixed number N of sets and then formulate a protocol that can use these sets, namely adding values to sets, removing values, or checking whether a value is contained in a set. The abstraction approach simply identifies all values that have the same memberships in these sets, yielding 2^N equivalence classes.

The limitation of AIF that our AIF-ω extension overcomes is that we can then only model a fixed number of honest agents: since if for instance each agent maintains its own set of public/private key-pairs, and the number of sets N needs to be fixed, so has to be the number of agents. Instead, AIF-ω allows for infinite families of sets, so one declares an infinite number of agents, and each has its own key ring. In practice this turns out to be very beneficial, even for a fixed number of agents: by avoiding the enumeration of agents and their sets in the set memberships, the verification for infinitely many agents is just as efficient as the verification for a single agent. This also greatly facilitates the verification of protocols of FutureID, as we can now verify security for an unbounded number of agents and servers, both of which an unbounded number can be honest and dishonest, respectively.

5 Establishing Methods and Languages for Privacy Goals

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 in Section 5.1, and the other to establish semantics for claims languages to allow reasoning over them. For the second sub-task, in Section 5.2, we define and unify the concepts and features of privacy-preserving attribute-based credentials (Privacy-ABCs), provide a language framework in XML schema, and give a formal semantics to describe the effects of the transactions in a privacy-friendly authentication system using Privacy-ABCs. Although not directly related with the goals of this task, we also describe in Section 5.3 a Prolog implementation for credential-policy matching.

5.1 Formal Methods for Privacy Goals

Previously on FutureID, we had developed the concept of α-β-privacy [265] (see also Deliverable D24.2) that gives a declarative interface to specifying privacy properties connected to standard low-level technical definitions of privacy. The idea is that one specifies merely the high-level information α that a system deliberately releases (e.g., in a zero-knowledge proof the statement being proved to the verifier) and from the specification of the system we obtain the low-level information β that the intruder (or any dishonest party) obtains from observing the network and its knowledge about the structure of the messages. Then α-β-privacy requires that the intruder cannot derive any high-level information from β that does not follow from α already.

We have applied α-β-privacy in a major case study to formally verifying the security and privacy properties of the FutureID architecture [263]. It turns out that α-β-privacy here gives us a *canonical* privacy goal: When showing a classical (non-zero knowledge-based) credential to a FutureID broker or to a service provider, then this provider can necessarily see all the attribute-value pairs in this credential. An honest server may for privacy delete this information, but we cannot prevent a dishonest—or honest but curious—service provider to somehow store this information permanently. Also we cannot prevent such a server from drawing conclusions from this information or to collaborate with other dishonest servers and pool their knowledge. However we can prevent that any further information is leaked. In this sense, α-β-privacy is the canonical privacy goal, when we set α to be exactly the information contained in the credential and β all the cryptographic messages available to the recipient.

Indeed our analysis shows that this holds as long as we can rely on the employed public-key infrastructure. Obviously, when the intruder can insert his own public keys into the system and get a user to accept them as the public key of a service provider or broker, then the intruder will see (but not be able to use) the credentials that the user intends to show to the actual server. In all other cases, privacy is preserved.

A relevant practical question is of course whether a server is indeed logging the credentials it has been shown. As argued before we cannot prevent a server from somehow secretly doing this, but there may be even a benign reason why a server wants to store this information: in case of a legal dispute a server could prove that it acted correctly by showing its logs. Such a log

can of course be a privacy risk, but in fact we can set up a system that lifts the server from this logging need without destroying accountability. We propose the idea of *onion-logging* which says that each server should encrypt the log information with a special public key (where the corresponding private key is stored in a particularly secured way); then the issued credential contains a hash of this encrypted log entry, which is passed along with the credential. The key is that the credential is not considered valid without a message that has the same hash as the message in the credential. When accusing a server for issuing a wrong credential, one would have to produce this message, and thus, the server, if honest, can decrypt this message and produce the corresponding log.

5.2 Concepts and Languages for Privacy-Preserving Attribute-Based Authentication

Privacy-preserving authentication mechanisms based on anonymous credentials, minimal disclosure tokens, self-blindable credentials, or group signatures [139, 77, 114, 117, 46, 335] offer a large variety of features. Similar features are often referred to by different names 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 several cryptographic protocols. This makes these technologies very difficult to understand, compare, and use.

We overcome these difficulties by providing 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 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 a cryptography-agnostic language framework and application programming interface (API) with well-documented data formats for credentials, 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. Our languages and API allow application developers to employ Privacy-ABCs without having to think about their cryptographic realization, similarly to how common cryptographic primitives such as encryption and signatures are used today: the application layer calls out to the cryptography through standardized interfaces; the concrete chosen algorithm is at most an initialization parameter. Our language has been implemented and will be made available as part of a reference implementation of a Privacy-ABC system which will include a number of cryptographic solutions.

Finally, we present a formal semantics that precisely defines the meaning of our comprehensive language and their expressed features. Such a rigorous mathematical description allows to determine, for instance, whether a user can fulfill a given authentication policy with her credential portfolio or whether a derived access token satisfies a policy. As our language covers the entire Privacy-ABC system, we also provide semantics that describe the intended system behaviour, i.e., the effects of state transitions— which are steered by our language—on the different entities and their knowledge states.

The life of a Privacy-ABC starts with an *issuance* process between a User and an Issuer. Our framework supports advanced issuance protocols where, for instance, attribute values or underlying credential secrets can be "carried over" from existing credentials or can be assigned jointly generated random values, without the Issuer being able to see the values. Credentials can subsequently be presented by Users to Verifiers to prove that certain requirements specified in a presentation policy are fulfilled. With Privacy-ABCs, Users can selectively reveal attributes from multiple credentials, or merely prove that the attributes satisfy a given predicate. Depending on the level of linkability that is desired, the User can thereby remain unlinkable and establish a fresh *pseudonym*—the privacy-friendly equivalent to a standard public key—, or re-authenticate under a previsouly used pseudonym. Credentials with *key binding* have an underlying secret key so that they can only be used in combination with that key. If the key is stored on a hardware device such as a smart card, the credential is essentially bound to the device. Multiple credentials and pseudonyms can be bound to the same secret key to securely tie several credentials together and to discourage credential sharing, i.e., Users lending their credentials to other Users.

Pseudonyms are the privacy-friendly analog to standard public keys, i.e., they are derived from a secret key and are given to a Verifier so that the User can later re-authenticate using the secret key. Unlike classical public keys, however, the User can generate an unlimited number of unlinkable pseudonyms from a single secret key.

To add accountability and prevent abuses, our framework further defines revocation and inspection of Privacy-ABCs. Regarding *revocation* we distinguish between two types: "issuer-driven" revocation where the Issuer renders a credential useless globally, and "verifier-driven" revocation where the Verifier can blacklist certain attribute values. *Inspection* is an optional mechanism that can be used to reveal user attributes in an encrypted form, so that they can only be recovered by a designated third party under well-specified circumstances, e.g., to allow anonymity lifting.

To enable applications to take full-advantage of the described features and mechanisms we subsequently propose data formats for the full life-cycle of Privacy-ABCs that express these features in a technology-agnostic way. Note that this is the first time that issuance, revocation, inspection and advanced concepts such as pseudonyms and key binding for Privacy-ABCs are modeled by an abstract language and are also fully supported by privacy-preserving policy and token formats. Our specification uses XML notation in the spirit of XML Schema and determines exact formats for the mechanism-independent information as well as anchor points for the opaque mechanism-specific cryptographic data. We also briefly describe the concept and functionality of an ABC-Engine which operates on top of the core cryptographic engine and contains all the mechanism-agnostic components of a Privacy-ABC system and processes and produces the data formats that we present.

Our work builds on the credential-based authentication requirements language (CARL) recently proposed by Camenisch et al. [120]. CARL allows a service provider (verifier) to specify which attributes a user needs to present, and by which issuer these attributes need to be certified, in order to get access. Compared to our work, CARL defines only a small part of a Privacy-ABC system, namely the presentation policy, but does not consider how these attributes are transmitted nor how credentials are issued or revoked. Bichsel et al. [59] have extended CARL

to cover the transmission of certified attributes. The current version of the U-Prove protocol natively provides a subset of features of our framework; the other features need to be added through extension points.

5.3 A Prolog Program for Matching Attribute-Based Credentials to Access Control Policies

In an attribute-based credential system [114, 117, 115, 46, 47, 104, 99, 235], users receive credentials from credentials issuers. A credential is a container of user attributes that are certified by the credential issuer. Users employ their credentials to be granted access to services that are protected by access control policies. An access control policy describes the credentials that a user must possess in order to be granted access to a service. An access control policy may describe the type of non-revoked credentials that a user must possess, the identity of the issuers of those credentials, the type of attributes that the credentials must contain and restrictions on the values of those attributes.

A user computes a presentation token in order to prove to the service provider that she possesses credentials that fulfill an access control policy. A presentation token consists of a description of the credentials information that the user reveals to the service provider in order to prove that her credentials fulfill the policy, and a cryptographic proof that guarantees that the user indeed possesses credentials with such information. This cryptographic proof certifies that the user credentials fulfill the policy, but does not disclose any other information on the user's credentials.

In order to compute a presentation token, a user must first check whether her credentials fulfill the access control policy. Some access control policies could be fulfilled by different subsets of the users credentials. For example, consider a policy that restricts access to the books offered by a library. The policy requires users to be members of the library, which they can prove if they possess a credential issued by the library, or if they are students and nationals of the country where the library is located, which they can prove if they possess an identity card and a student card that store the corresponding credentials. If a user possesses those three types of credentials, the user can choose which ones to use in order to compute the presentation token. This choice may have both efficiency and privacy implications: one the one hand, proving possession of the credential issued by library could be more efficient than proving possession of two credentials on an identity card and on a student card; on the other hand, if the library has few members, proving possession of the credential issued by the library hides the user identity only in a small set of users.

In the example above, it is easy to compute the different subsets of credentials that a user can employ to satisfy the policy. However, in general, the number of credentials a user possesses and the number of ways a policy can be satisfied can both be large. Additionally, presentation tokens can be associated to a pseudonym. Presentation tokens are in general unlinkable, i.e., the verifier does not know whether two tokens were computed by the same or by different users unless the policy that the tokens fulfill allows the verifier to link them. However, a policy may require tokens to be linked through a pseudonym, or may allow users to choose whether to link

her presentation tokens. Therefore, when computing a presentation token, a user is confronted with multiple combinations of credential subsets and pseudonyms.

We provide a Prolog program that, given a policy, allows the user to compute all the combinations of credentials and pseudonyms that can be employed to compute a presentation token that fulfills the policy. Prolog is a logic programming language and it is declarative, i.e., the program logic is expressed in terms of relations represented as facts and rules. Our Prolog program employs facts to represent the user pseudonyms and the user credentials information, such as the credential issuer, type and attributes, and employs rules to represent policies. By querying whether a rule is fulfilled by the existing facts, the Prolog engine computes and lists all the subsets of facts that fulfill the rule. Therefore, simply by representing users credentials as facts and policies as rules in Prolog, we obtain a program that outputs the desired credentials subsets. We remark thus the simplicity of our approach in comparison to using other programming paradigms, which would require the implementation of both a credential-policy matching algorithm to know whether a subset of credentials fulfill the policy, and an exhaustive search algorithm to list all the subsets of credentials that fulfill the policy.

Our Prolog approach is also useful for the verifier. The verifier's program represents the user pseudonyms and the user credential information disclosed by the user's presentation tokens as facts, while the policies are represented as rules. When the user wishes to access a new service, the verifier can check whether the user pseudonyms and credential information disclosed before already fulfill the access control policy associated to the new service. To do this, like in the user program, the verifier runs a query to check whether the rule that represents the policy is fulfilled by the existing facts. Thanks to this program, the verifier can spare the user from computing a new presentation token when the facts known by the verifier already fulfill the policy.

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

Task 24.3 conducts research on audits and data-handling mechanisms. For audits, in Section 6.1, we report on experiences during implementing blank digital signatures as well as optimizations that helped to improve their performance. Additionally, in Section 6.2, we propose a novel graph signature scheme, which makes it possible that an issuer certifies a committed graph, such that a prover can subsequently prove properties of the graph in zero-knowledge proofs of knowledge. Whereas this primitive is at the heart of the confidentiality-preserving topology certification mechanisms discussed in deliverable D24.2, we extend this work with a proof of the signature scheme's expressiveness. The signature scheme is capable of signing and proving knowledge of statements of NP languages, shown with a reduction to graph 3-colorability. Therefore, the signature scheme is a promising candidate for a wide range of applications.

For data-handling mechanisms, we focus our research on authentication mechanisms based on passwords and signatures and on privacy-preserving protocols that minimize the data that users have to disclose to service providers. We also work on existing eID solutions. In addition, we conduct research on computations on signed data and on data anonymization and data sharing between databases.

Password-Based Authentication. In Section 6.3, we describe a threshold password authenticated secret sharing protocol. In Section 6.4, we present two simple and extremely efficient proactively secure *distributed password verification* protocols.

Signature-Based Authentication. In Section 6.5, we give formal security definitions for a full-fledged privacy attribute-based credentials system. We provide a generic construction from lower-level building blocks that satisfies our definitions and we present secure instantiations of the building blocks. In Section 6.6, we propose a new kind of signature schemes, *unlinkable redactable block-signature* (URS) schemes, with which one can redact a signature and reveal only its relevant parts each time it is used. We construct an efficient URS scheme and we employ it to design the first universally composable anonymous credential system. It is also arguably one of the first such schemes to support efficient attribute disclosure with cost independent of the number of attributes in the issued credential without having to rely on random oracles. In Section 6.7 we present an efficient construction of round-optimal blind signature schemes in the standard model.

eID Solutions. In Section 6.8, we present design strategies for a privacy-friendly Austrian eID system in the public cloud.

Privacy-Preserving Protocols. In Section 6.9, we contribute a practical, secure and privacy-preserving mechanism enabling a Service Provider to verify whether the mobile phone of a given User currently resides within a certain geographical reference area at a given time. Our mechanism consists in having the location of the mobile phone determined by the Mobile Network Operator and certified using *anonymous credentials*. In Section 6.10, we propose a non-restricted and a restricted oblivious transfer with access control scheme.

Oblivious transfer with access control (OTAC) allows the sender to control access to the messages. The sender receives as input a list of messages and access control policies. Each receiver possesses a set of attributes, which is certified by a credential issuer, and is able to obtain a message only if the receiver's attributes satisfy the access control policy for that message. Receiver privacy requires that the sender does not get any information on the message obtained or on the receiver's attributes. In a non-restricted scheme, a receiver can obtain in one transfer phase all the messages whose access control policy is fulfilled by the receiver's attributes. In a restricted scheme, the receiver can only obtain one message per transfer phase. In Section 6.11, we revisit existing work on privacy-preserving billing. In privacy-preserving billing a meter measures a user's consumption of some utility or service and service providers apply fine-grained tariff policies, i.e., policies that require detailed and frequent consumption measurements, in order to determine the bill. Meters do not send consumption measurements to the service provider. Instead, the computation of the bill is done locally and only the amount to be paid is revealed to the service provider. We improve existing work in two ways. First, we generalize the security model to consider multiple meters and multiple users. Second, we propose a privacy-preserving billing protocol for our model that improves efficiency for policies described by splines.

Computations on signed data. In Section 6.12, we show how the service provider can perform computations on unencrypted signed data.

Data anonymization and sharing. In Section 6.13, we propose an (un)linkable pseudonym system to allow exchange of user data between databases. A converter serves as central hub to ensure controllability. The converter establishes individual pseudonyms for each server derived from a unique main identifier that every user has, but without learning the derived pseudonyms. The only information the converter still learns is that a server S_A wants to access data from a server S_B, which is the right amount of information to balance control and privacy.

6.1 Blank Digital Signatures: Optimization and Practical Experiences

In contrast to conventional digital signatures, involving a signer and a verifier, proxy-type digital signature schemes are signature schemes involving three parties, namely an originator, a proxy and a verifier. Here, the originator delegates the signing power (for some particular well defined set of messages) to a proxy. The proxy can then sign messages on behalf of the originator. Any verifier, given a message and a corresponding signature, can check whether the proxy has produced the signature on behalf of the originator (authenticity), the integrity of the message and whether the given message is one of the "allowed" messages.

Blank Digital Signatures (BDS) [206] are a special instance of proxy-type digital signatures, allowing an originator to define and issue a signature on a template, containing fixed and exchangeable elements. A designated proxy can then produce signatures for instantiations of this template (messages). More precisely, given a template signature, the proxy creates an instantiation by choosing one of the predefined values for each of the exchangeable elements and issues a signature with respect to the template signature. When verifying this signature, only the

message and the corresponding signature is needed, and it is required that the verifier does not learn anything about the unused choices in the exchangeable elements in the template (privacy property).

Blank Digital Signatures give rise to a lot of interesting applications, and, accordingly, the question arises how a BDS scheme would perform in a practical implementation, and to which extent it can be integrated into off-the-shelf cryptographic frameworks such as the Java Cryptography Architecture [289] and key infrastructures such as PKIX [154].

We propose optimizations for the BDS scheme in [206] and present a full-fledged implementation of this optimized version. Firstly, we briefly revisit the scheme and discuss possible practical applications. Then, we show how the scheme can be modified to use Type-3 pairings instead of the originally proposed Type-1 pairings and introduce optimizations for the encoding of templates. Subsequently, we show how the scheme can be integrated into the Java Cryptography Architecture and how the keying material can be encapsulated within X.509 certificates. Moreover, two possible signature formats, namely an XML and a PDF signature format, are proposed. Finally, timings of our implementation, showing the practical applicability of the BDS scheme, are provided and discussed.

6.2 Certification of Committed Graphs

Identity federation systems on a European scale will eventually interconnect a large number of independent subsystems, forming sprawling topologies. How these systems are interconnected yields security properties, such as which sub-systems depend on which other sub-systems or how information can flow in the overall system. We anticipate requirements to attest to the security of the system, largely integrity and availability properties, while keeping the blueprint of the identity federation system confidential. We have outlined in D24.2 how to certify topologies in such a way that the topology providers can prove in zero-knowledge proofs of knowledge security properties of the underlying infrastructure [196].

As solution to the problem of topology certification, we propose a graph signature and proof system [197] that is capable of issuing digital signatures on graph representations, such that a prover can access the elements of the graph in zero-knowledge proofs of knowledge. The key idea for this method is to encode the graph in a structured Gödel Numbering, such vertices and labels are represented as prime numbers and their combinations as prime products. This idea draws inspiration from efficient attributes for anonymous credential systems and the Camenisch-Groß encoding [100]. Given such a Gödel Numbering the graph is embedded into the Camenisch-Lysyanskaya signature scheme [116], such that the graph elements can still be accessed with divisibility and co-primality proofs.

The foundational construction [197] investigates the core idea for the signature scheme and evaluates its expressivity. To that end, it is proven that one can represent the Graph 3-Colorability problem in the graph signature scheme. As Graph 3-Colorability is NP-complete, one can reduce statements from any NP languages into a form that can be signed by the graph signature scheme. Therefore, the graph signature scheme is similarly expressive as zero-knowledge proofs of knowledge. This shows that it is feasible in principle to have signature schemes on statements

from NP languages that are enabled to have zero-knowledge proofs of knowledge on them.

The proposal for topology certification defined proof predicates and constructions for connectivity and isolation proofs [196]. The isolation proofs are of particular interest, as they set this scheme apart from other proposals in transitive or homomorphic signatures. The problem here is that a prover might choose to "forget" edges or vertices to prove isolation. Hence, a graph signature scheme needs to bind all elements of the graph together such that the prover cannot evade mentioning a vertex or an edge. It still makes it possible to prove isolation efficiently by determining a bipartition of the edge set which yields the isolation statement, computing the cumulative product of the two partitions in a commitment and subsequently proving that both products are co-prime.

Related Work The proposed primitive of graph signatures relates to the concept of transitive signatures [260]. In a transitive signature scheme, the signatures of edges can be combined in such a way that the resulting signature is a valid signature on the path of the edges in question. This has the advantage that edges can be signed separately and combined to signatures on paths thereafter. The existing schemes in this field, however, do not allow for zero-knowledge proofs of knowledge on complex graph properties, such as isolation.

Our Contribution The key contribution made in FutureID is a novel signature scheme and corresponding zero-knowledge proof system that allows certification of graph representations [197]. The scheme is the first proposal to certify arbitrary undirected vertex- and edge-labeled graphs and allow full access to all components of the graph in subsequent proofs. The scheme allows for a joint issuing of signatures on compositions of issuer-known sub-graphs and user-generated committed/hidden sub-graphs. Thereby, it allows to bootstrap new graph signatures from existing graph signatures.

The initial construction of the signature scheme comes with a reduction proof showing that statements from an NP-complete language can be embedded in the signature scheme.

Furthermore, more practical constructions for topology certification [196] contribute proofs for connectivity (i.e., that there exists a connected path of length ℓ) and isolation (i.e., that there is no path between two specified vertices).

Overall, this work shows feasibility of graph signatures and zero-knowledge proofs thereon as a new area in cryptography. The graph signatures can be used in a variety of situations, either to certify topologies, such as virtualized infrastructures or identity federation systems, or to certify and prove properties of graph representations, such as provenance graphs.

6.3 Threshold password-authenticated secret sharing

Properly protecting our digital assets still is a major challenge today. Because of their convenience, we protect access to our data almost exclusively by passwords, despite their inherent weaknesses. Indeed, not a month goes by without the announcement of another major password

breach in the press. In 2013, hundreds of millions of passwords were stolen through server compromises, including massive breaches at Adobe, Evernote, LivingSocial, and Cupid Media. In August 2014, more than one billion passwords from more than 400,000 websites were reported stolen by a single crime ring. Barring some technical blunders on the part of Adobe, most of these passwords were properly salted and hashed. But even the theft of password hashes is detrimental to the security of a system. Indeed, the combination of weak human-memorizable passwords (NIST estimates sixteen-character passwords to contain only 30 bits of entropy [87]) and the blazing efficiency of brute-force dictionary attacks (currently testing up to 350 billion guesses per second on a rig of 25 GPUs [194]) mean that any password of which a hash was leaked should be considered cracked.

Stronger password hash functions [300] only give a linear security improvement, in the sense that the required effort from the attacker increases at most with the same factor as the honest server is willing to spend on password verification. Since computing password hashes is the attacker's core business, but only a marginal activity to a respectable web server, the former probably has the better hardware and software for the job.

A much better approach to password-based authentication, first suggested in [179], is to distribute the capability to test passwords over multiple servers. The idea is that no single server by itself stores enough information to allow it to test whether a password is correct and therefore to allow an attacker to mount an offline dictionary attack after having stolen the information. Rather, each server stores an information-theoretic share of the password and engages in a cryptographic protocol with the user and the other servers to test password correctness. As long as less than a certain threshold of servers are compromised, the password and the stored data remain secure.

Building on this approach, several threshold password-authenticated key exchange (TPAKE) have since appeared in the literature [179, 220, 253, 74, 166, 327, 229], where, if the password is correct, the user shares a different secret key with each of the servers after the protocol. Finally addressing the problem of protecting user data, threshold password-authenticated secret sharing (TPASS) protocols [37, 118] combine data protection and user authentication into a single protocol. They enable the password-authenticated user to reconstruct a strong secret, which can then be used for further cryptographic purposes, e.g., decrypting encrypted data stored in the cloud. An implementation of the protocol by Brainard et al. [74] is commercially available as EMC's *RSA Distributed Credential Protection* (DCP) [171].

Unfortunately, all protocols proposed to date do not provide satisfying security. Indeed, for protocols that are meant to resist server compromise, their authors are surprisingly silent about what needs to be done when a server actually gets corrupted and how to recover from that. The work by Di Raimondo and Gennaro [166] is the only one to mention the possibility to extend their protocol to provide proactive security by refreshing the shares between time periods; unfortunately, no details are provided. The RSA DCP product description [171] mentions a re-randomization feature that "can happen proactively on an automatic schedule or reactively, making information taken from one server useless in the event of a detected breach." This feature is not described in any of the underlying research papers [74, 327], however, and neither is a security proof known. Taking only protocols with provable security guarantees into account,

the existing ones can protect against servers that are malicious from the beginning, but do not offer any guarantees against adaptive corruptions. The latter is a much more realistic setting, modelling for instance servers getting compromised by malicious hackers. This state of affairs is rather troubling, given that the main threats to password security today, and arguably, the whole *raison d'être* of TPAKE/TPASS schemes, come from the latter type of attacks.

One would hope to be able to strengthen existing protocols with ideas from proactive secret sharing [212] to obtain security against adaptive corruptions, but this task is not straightforward and so far neither the resulting protocol details nor the envisaged security properties have ever been spelled out. Indeed, designing cryptographic protocols secure against adaptive corruptions is much more difficult than against static corruptions. One difficulty thereby is that in the security proof the simulator must generate network traffic for honest parties *without* knowing their inputs, but, once the party is corrupted, must be able to produce realistic state information that is consistent with the now revealed actual inputs as well as the previously simulated network traffic. Generic multiparty computation protocols secure against adaptive corruption can be applied, but these are too inefficient. In fact, evaluating a single multiplication gate in the most efficient two-party computation protocol secure against adaptive corruptions [98] is more than three times slower than a full execution of the dedicated protocol we present here.

We provide the first threshold password-authenticated secret sharing protocol that is provably secure against *adaptive* corruptions, assuming data can be securely erased, which in this setting is a standard and also realistic assumption. Our protocol is a two-server protocol in the public-key setting, meaning that servers have trusted public keys, but users do not. We do not require random oracles. We also describe a *recovery procedure* that servers can go execute to recover from corruption and to renew their keys assuming a trusted backup is available. The security of the password and the stored secret is preserved as long as both servers are never corrupted simultaneously.

We prove our protocol secure in the universal composability (UC) framework [128]. The very relevant advantages of composable security notions for the particular case of password-based protocols have been argued before [131, 118]; we briefly summarize them here. In composable notions, the passwords for honest users, as well as their password attempts, are provided by the environment. Passwords and password attempts can therefore be distributed arbitrarily and even dependently, reflecting real users who may choose the same or similar passwords for different accounts. It also correctly models typos made by honest users when entering their passwords: all property-based notions in the literature limit the adversary to seeing transcripts of honest users authenticating with their correct password, so in principle security breaks down as soon as a user mistypes the password. Finally, composable definitions absorb the inherent polynomial success probability of the adversary into the functionality. Thus security is retained when the protocol is composed. In contrast, composition of property-based notions with non-negligible success probabilities is problematic because the adversary's advantage may be inflated. Also, strictly speaking, the security provided by property-based notions is guaranteed only if a protocol is used in isolation.

Our construction uses the same basic approach as the TPASS protocols of Brainard et al. [74] and Camenisch et al. [118]. During Setup, the user generates shares of his key and password

and sends them to the servers (together with some commitments that will later be used in Retrieve). During Retrieve, the servers run a subprotocol with the user to verify the latter's password attempt using the commitments and shares obtained in Setup. If the verification succeeds, the servers send the shares of the key back to the user, who can then reconstruct the key. Furthermore, the correctness of all values exchanged is enforced by zero-knowledge proofs. Like the recent work of Camenisch et al. [111], we do not require the user to share the password during Retrieve but run a dedicated protocol to verify whether the provided password equals the priorly shared one. This offers additional protection for the user's password in case he mistakenly tries to recover his secret from servers different from the ones he initially shared his secret with. During setup, the user can be expected to carefully choose his servers, but retrieval happens more frequently and possibly from different devices, leaving more room for error.

The novelty of our protocol lies in how we transform the basic approach into an efficient protocol secure against an adaptive adversary. The crux here is that parties should never be committed to their inputs but at the same time must prove that they perform their computation correctly. We believe that the techniques we use in our protocol to achieve this are of independent interest when building other protocols that are UC-secure against adaptive corruptions. First, instead of using (binding) encryptions to transmit integers between parties, we use a variant of Beaver and Haber's non-committing encryption based on one-time pads (OTP) [45]: the sender first commits to a value with a mixed trapdoor commitment scheme [98] and then encrypts both the value and the opening with the OTP. This enables the recipient to later prove statements about the encrypted value. Second, our three-party password-checking protocol achieves efficiency by transforming commitments with shared opening information into an Elgamal-like encryption of the same value under a shared secret key. To be able to simulate the servers' state if they get corrupted during the protocol execution, each pair of parties needs to temporarily re-encrypt the ciphertext with a key shared between them.

Finally, we note that our protocol is practical: users and servers have to perform a few hundred exponentiations each, which translates to an overall computation time of less than 0.1 seconds per party.

6.4 Optimal Distributed Password Verification

We present two simple and extremely efficient proactively secure *distributed password verification* protocols, allowing a login server $\mathcal{LS}$ and a number of back-end servers $\mathcal{S}_1, \ldots, \mathcal{S}_n$ to jointly determine the correctness of a user's password, while ruling out offline dictionary attacks unless *all* servers are corrupted during the *same* time period. A corrupt $\mathcal{LS}$ only sees the passwords of user accounts that are created or logged into during the corruption. No passwords, password hashes, or any other offline-attackable information is leaked for accounts that are inactive during the corruption. We think this is a reasonable compromise for not requiring user-side software, as it provides adequate protection against "smash-and-grab" attacks and short-term corruptions.

Login, i.e., password verification, is a single-round protocol requiring just one exponentiation in a prime-order group on each server (two for $\mathcal{LS}$), which is essentially optimal unless schemes without public-key operations can be found. The recovery and key refresh procedure is non-

interactive and only involves a couple of additions and pseudo-random function evaluations per server, making it more than efficient enough to perform it preventively on a regular basis instead of just after a detected breach. Our first construction works in any prime-order group, including elliptic curves, and involves a three-round account creation (password setup) protocol with three exponentiations per server (six for $\mathcal{LS}$). Our second construction is based on elliptic curves with bilinear maps and also offers single-round account creation with one exponentiation per back-end server and one exponentiation and one pairing computation for $\mathcal{LS}$. Both our protocols assume that the key refresh procedure has access to a special backup tape that is not connected during normal operation. In practice, this can be achieved by using smart cards or by making use of properties of modern cloud platforms, as we will explain.

Given their extreme efficiency, it is all the more surprising that we managed to prove our constructions secure under a very strong universally composable (UC) [128] notion with transient corruptions. Parties can be dynamically corrupted at any point in the protocol, even between communication rounds. Transiently corrupted parties leak their full state, but not the content of their backup tape, to the adversary and remain corrupted until the next key refresh. Permanently corrupted parties additionally leak the backup tape and cannot be recovered.

As was argued before [229, 118, 111, 97], universal composability offers important advantages over traditional game-based definitions in the particular case of password-based protocols. Namely, UC notions leave the choice of passwords to the environment, so that arbitrary distributions and dependencies between passwords are correctly modeled. This is crucial to guarantee security in real-life settings where users make typos when entering their passwords, share passwords, or use the same password for different accounts—none of which are covered by currently known game-based notions. Also, it is very unclear whether protocols can be securely composed with the non-negligible attack probabilities that game-based definitions tend to employ. We prove our constructions secure in the random-oracle model under the (gap) one-more Diffie-Hellman assumption that was previously used to prove security for blind signature [64], oblivious transfer [147], TPASS protocols [221], and set intersection protocols [222].

We achieved this rare combination of strong security and high efficiency by careful proof techniques in the random-oracle model, as well as through some of compromises in security that are very reasonable for practical use, but save on cryptographic machinery in the protocol design. First, we assume that the initialization of all servers takes place in a trusted environment where all servers are honest. During initialization, we assume that $\mathcal{LS}$ can transmit one secure message to each back-end server $\mathcal{S}_i$. This secure initialization is not hard to achieve in practice. Server refresh, i.e., whereby a server can recover from a transient corruption, does not require any interaction with other servers.

Second, the back-end servers $\mathcal{S}_1, \ldots, \mathcal{S}_n$ do not learn which user is logging in or whether the password was correct. This definitely limits their ability to throttle failed login attempts, but since $\mathcal{LS}$ can apply clever throttle algorithms based on user id and login results, the natural throttling of back-end servers just by requiring network communication should suffice to fend off attacks. Finally, we do not cover robustness: an adversary can make $\mathcal{LS}$ "err on the safe side" and conclude that the password was false while in fact it was correct—but not the other way around. This could be fixed by adding the same zero-knowledge or pairing verification as during

account registration. This would have a major impact on efficiency, however, so we prefer to accept this rather benign attack in the model.

As a technical contribution, our scheme employs a novel technique to obtain proactive security that may be of independent interest. In a nutshell, we start off from a basic scheme that is secure under dynamic but non-transient corruptions. The basic scheme is secure under the gap one-more Diffie-Hellman assumption, but the security proof requires guessing one server at the beginning of the game that will not get corrupted during the game. This guessing induces a tightness loss in the reduction equal to the number of servers. While that loss could still be tolerated, things get worse when moving this scheme into a proactive setting. Here one would have to guess an uncorrupted server at the beginning of each epoch, so that the tightness loss blows up exponentially in the number of epochs. An easy but unsatisfying solution could be to restrict the scheme to a logarithmic number of epochs, or to only model semi-static corruptions where the adversary has to announce all servers that it wants to corrupt at the beginning of each epoch. Instead, we modify the scheme to apply random-oracle-generated blinding factors to all protocol messages, so that protocol messages do not commit servers to their keys, without ruining the overall functioning of the protocol. In the simulation, we can therefore choose a server's keys only at the moment that it is corrupted and carefully program the random oracle to ensure consistency of previous protocol messages, without having to guess anything upfront.

Of the threshold password-authenticated protocols in the literature, only Camenisch et al. [97] describe a recovery procedure and prove their protocol secure against transient corruptions. Proactive security in the protocol of Camenisch et al. [97] unfortunately comes at a considerable cost: "a few hundred exponentiations" per server may be within practical reach for occasional data retrieval, but not for high-volume password verification.

6.5 Formal Treatment of Privacy-Enhancing Credential Systems

Privacy-enhancing attribute-based credentials systems (aka PABCs, *anonymous credentials*, or *pseudonym systems*) allow for cryptographically strong user authentication while preserving the users' privacy by giving users full control over the information they reveal. There are three types of parties in a PABC system. *Issuers* assign sets of attribute values to *users* by issuing credentials for these sets. Users can present (i.e., prove possession of) their credentials to *verifiers* by revealing a subset of the attributes from one or more credentials. The verifiers can then check the validity of such presentations using the issuers' public keys, but they do not learn any information about the hidden attributes and cannot link different presentations by the same user. This basic functionality of a PABC system can be extended in a large number of ways, including pseudonyms, revocation of credentials, inspection, proving relations among attributes hidden in presented credentials, and key binding [108, 93].

The importance of privacy and data minimization in authentication systems has been emphasized, e.g., by the European Commission in the European privacy standards [6, 3] and by the US government in the National Strategy for Trusted Identities in Cyberspace (NSTIC) [217]. With IBM's Identity Mixer based on CL-signatures [124, 114, 116, 117] and Microsoft's U-Prove based on Brands' signatures [77, 293], practical solutions for PABCs exist and are cur-

rently deployed in several pilot projects [108, 4, 218, 5]. In fact, numerous anonymous credential schemes as well as special cases thereof such as group signatures, direct anonymous attestation, or identity escrow have been proposed, offering a large variety of different features [139, 77, 117, 335, 114, 46, 187, 138, 293, 279, 346, 40].

Despite this large body of work, a unified definitional framework for the security properties of a full-fledged PABC system is still missing. Existing schemes either have targeted security definitions for specific use cases [114, 46, 187, 138] or do not provide provable security guarantees at all [293, 346, 279]. One possible reason for the lack of a generic framework is that dedicated schemes for specific scenarios are often more efficient than generic solutions. However, defining, instantiating, and re-proving tailored variants of PABCs is hard and error-prone. Clearly, it is desirable to have a unified definitional approach providing security definitions for a full-fledged PABC system. It turns out that achieving such definitions is far from trivial as they quickly become very complex, in particular, if one allows for relations between hidden attributes when issuing and presenting credentials, as we shall discuss. Nevertheless, we take a major step towards such a unified framework for PABCs by formally defining the most relevant features, detached from specific instantiations or use cases. We further provide a generic construction of a PABC system based on a number of simpler building blocks such as blind signatures or revocation schemes, and a formal proof that this construction meets our security definitions. Finally, we give concrete instantiations of these components.

Considered Features. Our definition of PABC systems comprises the richest set of features integrated into a holistic PABC scheme so far. It supports credentials with any fixed number of attributes, of which any subset can be revealed during presentation. A single presentation can reveal attributes from *multiple credentials*. Users can *prove equality* of attributes, potentially across different credentials, without revealing their exact values. Users have *secret keys* from which arbitrarily many *scope-exclusive pseudonyms* can be derived. That is, for a given secret key and *scope*, only one unique pseudonym can be derived. Thus, by reusing the same scope in multiple presentations, users can intentionally create linkability between presentations; using different scopes yields mutually unlinkable pseudonyms.

Credentials can optionally be bound to the users' secret keys to prevent users from sharing their credentials. When performing a presentation that involves multiple credentials and/or a pseudonym, all credentials and the pseudonym must be bound to or derived from the same user secret key, respectively. Issuers can *revoke credentials*, so that they can no longer be used in presentations. During issuance, some attribute values may be hidden from the issuer or *"carried over"* from existing credentials. This latter *advanced issuance* means that the issuer does not learn their values but is guaranteed that they are equal to an attribute in an existing credential.

Our Contributions. We give formal security definitions for a full PABC system that incorporates all of the features mentioned above. We provide a generic construction from lower-level building blocks that satisfies our definitions and we present secure instantiations of the building blocks.

In terms of security, informally, we expect presentations to be *unforgeable*, i.e., users can only

present attributes from legitimately obtained and unrevoked credentials, and to be *private*, i.e., they do not reveal anything more than intended. For privacy, we distinguish weak privacy, where presentations of a credential cannot be linked to a specific issuance session, and the strictly stronger notion of simulatable privacy, where in addition presentations of the same credential cannot be linked to each other. This allows us to cover (slight variants of) the most prevalent schemes used in practice, U-Prove and Identity Mixer.

Formally defining these properties is far from trivial because of the complexity of our envisaged system. For example, user can obtain credentials on (*i*) revealed, (*ii*) blindly carried-over, and (*iii*) fully blind attributes. Each type comes with different security expectations that must be covered by a single definition. Carried-over attributes, for example, present a challenge when defining unforgeability: While the issuer never learns the attributes, the adversary must not be able to present a value that was not previously issued as part of a pre-existing credential. For privacy, one challenge is to formalize the exact information that users intend to reveal, as they might reveal the same and possibly identifying attributes in different presentations. Revocation gives the issuer the power to exclude certain credentials from being used, which must be modeled without cementing trivial linking attacks into the model that would turn the definition moot.

The resulting definitions are rather complex, and thus proving that a concrete scheme satisfies them from scratch can be a challenging and tedious task. Also, proofs for such monolithic definitions tend to be hard to verify. We therefore define a set of building blocks, strongly inspired by existing work, and show how they can be generically composed to build a secure PABC system. Our construction is efficient in the sense that its complexity is roughly the sum of the complexities of the building blocks. Additionally, this construction allows for simple changes of individual components (e.g., the underlying revocation or pseudonym schemes) without affecting any other building blocks and without having to reprove the security of the system. Finally, we give concrete instantiations for all our building blocks based on existing protocols.

Related Work. Our definitions are inspired by the work of Chase et al. [135, 46], who provide formal, property-based definitions of delegatable anonymous credential systems and give a generic construction from so-called P-signatures [47]. However, their work focuses on pseudonymous access control with delegation, but lacks additional features such as attributes, revocation, and advanced issuance.

PABCs were first envisioned by Chaum [139, 142], and they have been a vivid area of research over the last decade. The currently most mature solutions are IBM's Identity Mixer based on CL-signatures [124, 114, 116, 117] and Microsoft's U-Prove based on Brands' signatures [77, 293]. A first formal definition by Camenisch and Lysyanskaya [114] in the ideal-real world paradigm covered the basic functionalities without attributes and revocation. Their definition is standalone and does not allow composability as honest parties never output any cryptographic values such as a credentials or pseudonyms. This restriction makes it infeasible to use their schemes as building block in a larger system. These drawbacks are shared by the definition of Garman et al. [187].

A recent MAC-based credential scheme [138] allows for multiple attributes per credential, but requires issuer and verifier to be the same entity. Furthermore, it does not cover pseudonyms,

carry-over attributes, or revocation and provides rather informal security definitions. Similarly, the recent work of Hanser and Slamanig [208] neither considers any of these features nor blind issuance, i.e., the issuer always learns all the user's attributes.

Baldimtsi and Lysyanskaya [39] define a blind signature scheme with attributes and claim that it yields an efficient *linkable* anonymous credential scheme, again without giving formal definitions. The scheme can be seen as a weakened version of our signature building block without unlinkability or extractability of hidden attributes – properties that are crucial for our PABC system.

Finally, existing definitions of attribute-based signatures, e.g., [241, 254, 317] differ substantially from those of our building block for privacy-enhancing attribute-based signature, as again they do not consider, e.g., blind issuance.

6.6 Unlinkable Redactable Signatures and Their Applications to Anonymous Credentials

Digital signature schemes are a fundamental cryptographic primitive. Besides their use for signing digital documents and building classical cryptographic protocols such as key exchange, they are also important components of more advanced cryptographic protocols, such as blind signatures [177, 18], group signatures [50, 54, 230], direct anonymous attestation [80], e-cash [143], voting schemes [214], adaptive oblivious transfer [122, 95] and anonymous credentials [47].

The efficient construction of such protocols, however, demands special properties of a signature scheme, in particular, when the protocol needs to achieve strong privacy properties. The most important such properties seem to be the issuance of a signature and its use later in a protocol in an *unlinkable* way as well as the scheme being able to sign *blocks* of messages. The first signature scheme that fulfilled these requirements is a blind signature scheme by Brands [75]. The drawback of blind signature schemes, however, is that if the user uses a signature later on in the designed protocol, typically to convince a third party that she has obtained a signature from the signer on some message, she needs to reveal the actual signature so that the third party can verify the authenticity of the signature. Thus, such a signature can typically be used only once, which turns out to be quite limiting for applications like group signatures, multi-show anonymous credentials, and compact e-cash [101].

Camenisch and Lysyanskaya [116, 117] were the first to propose signature schemes (CL-signatures) overcoming these drawbacks. Their schemes are secure under the Strong RSA or the q-SDH assumption, respectively, and allow for an alternative approach when using a signature in a protocol: instead of revealing the signature to a third party, the user can employ zero-knowledge proofs to convince the third party that she possesses a valid signature from the signer. While in theory this is possible for any signature scheme, CL-signatures were the first that allowed one to do this efficiently with so-called generalized Schnorr proofs of knowledge of and about discrete logarithms. This is achieved because of the algebraic properties of CL-signatures, in particular, because no hash function is applied to the message and the signature and message values are either exponents or group elements.

With the advent of CL-signatures, the area of privacy preserving protocols flourished considerably and numerous new protocols based on them have been proposed. This has also made it apparent, that CL-signatures still have a number of drawbacks:

(1) In order to make generalized Schnorr proofs of knowledge non-interactive (which is often required in applications), one needs to resort to the Fiat-Shamir heuristic, i.e., the random oracle model and thus looses all provable security guarantees for the concrete scheme where the random oracle is instantiated by a hash function [129].

(2) When designing a protocol to be secure in a Universal Composability model [128], where no rewinding can be used to prove security, witnesses in the generalized Schnorr proofs of knowledge need to be encrypted under a public key that is a part of a common reference string (CRS). As the witnesses (messages signed with CL-signatures) are discrete logarithms, such encryption is rather expensive [41, 103] and most often renders the overall protocol unpractical.

(3) When proving ownership of a CL-signature on many messages, all the messages are needed for the verification of the signature and therefore the proof of possession of a signature will be linear in the number of messages.

Following the work of Groth and Sahai [198] which allowed for the first time to construct efficient non-interactive zero-knowledge proofs (NIZK) without using random oracles, albeit for a limited set of languages, many works focused on so-called structure-preserving signatures [14, 15, 13]. These signature schemes allow seamless integration with Groth-Sahai proofs and can also be used as signatures of knowledge [137], hence applicable in the scenarios in which one could previously use CL-signatures.

While overcoming the first two shortcomings listed above, structure-preserving signatures also suffer in terms of performance when signing multi-block messages, which is typically the case for applications such as anonymous credentials. Indeed, like in the case of CL-signatures, the size of proofs of structure-preserving signatures grows linearly with the number of messages. As the constant factor is larger than the one for generalized Schnorr, they lose their attractiveness as building block for applications requiring multi-block messages. Moreover, it seems this shortcoming cannot be overcome for fundamental reasons as a Groth-Sahai proof-of-knowledge of a signature on hidden group elements can be seen as a degenerate commitment, but group to group commitments do not shrink [17].

Our contribution. We take a different approach that has the goal of both being employable by real-world applications and addressing the above drawbacks: we propose a new kind of signature schemes, *unlinkable redactable block-signature* (URS) scheme with which one can redact a signature and reveal only its relevant parts each time it is used. Moreover, URS are unlinkable and the same signature can be redacted and revealed multiple times without linking two redacted signatures which were derived from the same signature. We view our contribution as threefold:

First, we formally defined URS schemes. We present property-based security definitions for URS schemes, i.e., *unlinkability* and *unforgeability*. To validate that these definitions are sufficient,

we provide a UC functionality for unlinkable redactable signatures and then show that a secure URS scheme implements the UC functionality.

Second, we construct an efficient URS scheme from vector commitments [228, 243, 132], structure-preserving signatures [14, 15], and (a minimal dose of) non-interactive proofs of knowledge (NIPoK), which in practice can be instantiated by Groth-Sahai proofs [198]. As we are interested in practical efficiency, we deliberately instantiate our scheme with concrete building blocks, and rely on stronger assumptions. Our modular approach does, however, facilitate the construction of 'paranoid' instantiations, e.g., by falling back on the CDH based vector commitment scheme of [132], expectedly at a large penalty in performance. We show how to make use of algebraic properties in our building blocks to minimize the witness size of the NIPoK. Our main target is Groth-Sahai proofs, but we also discuss extensions for optimizing the efficiency of generalized Schnorr proofs of knowledge in case one is willing to accept random oracles.

Third, to demonstrate the versatility of our URS scheme as a CL-signature 'replacement', we employ it to design the first universally composable anonymous credential system. It is also arguably one of the first such schemes to support efficient attribute disclosure with cost independent of the number of attribute in the issued credential without having to rely on random oracles. Even in the random oracle model this has so far been an elusive goal. The state of the art, [99], supported only a constant number of attributes, or would otherwise require an artificially large RSA modulus for the credential scheme.

More precisely, to construct anonymous credentials we require a blind issuing protocol (that allows users to obtain credentials using pseudonyms) and a protocol to prove possession of credentials for the revealed attributes. While the latter follows for free from unlinkable redactable signatures, the former needs a specific protocol. Because of efficient selective disclose, our scheme is very attractive and quickly surpasses schemes based on CL-sigantures or blind signatures [76, 117] when the number of attributes grows.

Other immediate applications of our URS scheme include efficient e-cash, credential-based key exchange, e-voting, and others.

Related work. A variety of signature schemes with flexible signing capabilities and strong privacy properties have been proposed [136, 73, 36, 49]. Such schemes generalize existing privacy preserving schemes such as group signatures, and anonymous credentials. Our work is most closely related to the work of Chase et al. [136] on controlled malleable signatures. They define simulatable malleable signatures and show how to instantiate (a base scheme without attributes for) delegatable credentials.

In principle, malleable signatures allow the owner of a signature to hide a part of the message, while still being able to produce a signature that verifies for the remaining message. Anonymity is implied by the context hiding property of malleable signatures, which requires derived signatures to be indistinguishable from freshly generated signatures on the derived messages. Existing malleable signature constructions, however, do not meet our requirements in terms of efficiency. Ahn et al. [24] and Attrapadung et al. [34] consider redactable signatures but for a different redact operation that allows to quote substrings. While these works do not aim at being directly

applicable in real world protocols, especially when redacting a block message with a large number of attributes, they provide a fresh definitional approach. Indeed, our definitions of unlinkability and unforgeability are inspired by [136], which in turn build on [24].

Vector commitments [133, 243, 132], a cryptographic building block introduced by Catalano, Fiore and Messina [133], are our starting point for achieving practical efficiency. Kate et al. [228] suggest how to build one-time-show credentials by using their polynomial commitments scheme to commit to the attributes, which the issuer then certifies by signing the resulting commitment. While not supporting multi-show, the work by Kate et al. is a useful generic recipe and our work can be seen as lifting their intuition to multi-show credentials and putting it on a formal basis. Also, recently, Kohlweiss and Rial [233] built a private access control protocol with efficient selective attribute disclosure from vector commitments. Their solution, however, requires interaction.

The first efficient multi-show anonymous credential scheme is [114]. Its success led to implementations [124] and related protocols are even contained in standards [80]. Many cryptographic papers on anonymous credentials often focus on a base protocol, with pseudonyms, but without attributes. A notable exception is [99], which studies an extension to the identity mixer system for efficient attribute disclosure. Their result is in the random oracle model, however.

The first non-interactive anonymous credential scheme that does not rely on random oracles is [47]. P-signatures, the primitive underlying their construction is a signature scheme falling short of being structure-preserving and is engineered towards the anonymous credentials application. Anonymous delegatable credentials [46, 182] allow the owner of a credential to not only prove access rights, but also to delegate them. Both P-signatures, and the work on delegatable credentials consider a base protocol without attributes. Belenkiy et al. [48] define multi-block P-signatures with a proof size linear to the number of blocks.

Extending the work of [47] in combination with vector commitments, Izabachène et al. [219] constructed block-wise P-signatures and built efficient anonymous credentials out of it. Although their techniques are similar to ours, we list some crucial differences here. First, we aim for modularity of constructions and proofs, formalizing security properties for all building blocks, including vector commitments, from which we build our URS and then anonymous credentials. Thus, all building blocks we define can be used for modular design of other protocols (e-cash, group signatures, e-voting, etc.). Furthermore, we provide UC definitions that allow one to build UC-secure protocols based on our URS scheme. Second, their credential scheme targets a restricted class of predicates, individual attribute release and dot-products, and for revealing multiple attributes is less (at least twice) efficient than ours and requires more communication rounds. Instead we advocate the use of general-purpose NIZK to prove arbitrary predicates as natural extensions to our core scheme. Third, there is no prevention of a replay attack for presentation (showing the credentials) algorithm in [219] - the adversary can reuse the proof from another user and this is not considered a forgery. Indeed their definitions do not guarantee that credentials are non-malleable, implying that multiple credential proofs of true predicates may be combinable into a proof of an as of yet unproven, but true, predicate. This is crucial for a non-interactive credential scheme and is taken care of in this work.

Canard and Lescuyer [127] proposed an anonymous credential system built from sanitizable

signatures—a special case of malleable signatures. Apart from being less efficient than ours, their scheme is only proven secure in the random oracle model.

6.7 Practical Round-Optimal Blind Signatures in the Standard Model

The concept of blind signatures [140] dates back to the beginning of the 1980s. A blind signature scheme is an interactive protocol, where a user (or obtainer) requests a signature on a message which the signer (or issuer) must not learn. In particular, the signer must not be able to link a signature to the execution of the issuing protocol in which it was produced (*blindness*). Furthermore, it should even for adaptive adversaries be infeasible to produce a valid blind signature without the signing key (*unforgeability*). Blind signatures have proven to be an important building block for cryptographic protocols, most prominently for e-cash, e-voting and one-show anonymous credentials. In more than 30 years of research, many different (> 50) blind signature schemes have been proposed. The spectrum ranges from RSA-based (e.g., [140, 106]) over DL-based (e.g., [287, 10]) and pairing-based (e.g., [64, 62]) to lattice-based (e.g., [310]) and code-based (e.g., [292]) constructions as well as constructions from general assumptions (e.g., [9, 209, 177]).

Blind signatures and their round complexity. Two distinguishing features of blind signatures are whether they assume a common reference string (CRS) set up by a trusted party to which everyone has access; and the number of rounds in the signing protocol. Those which require only one round of interaction (two moves) are called *round-optimal* [177]. Besides improving efficiency, round optimality also directly yields concurrent security (which otherwise has to be dealt with explicitly; e.g., [231, 209]). There are very efficient round-optimal schemes [141, 64, 51] under interactive assumptions (chosen-target-one-more-RSA-inversion and chosen-target-CDH assumption, respectively) in the random oracle model (ROM), as well as under the interactive LRSW [250] assumption in the CRS model [190]. All these schemes are in the honest-key model, where blindness only holds against signers whose keys are generated by the experiment.

Fischlin [177] proposed a generic framework for constructing round-optimal blind signatures in the CRS model with blindness under malicious keys: the signer signs a commitment to the message and the blind signature is a NIZK proof of a signed commitment which opens to the message.

Using structure-preserving signatures [14] and the Groth-Sahai (GS) proof system [198] instead of general NIZKs, this framework was efficiently instantiated in [14]. In [62, 63], Blazy et al. gave alternative approaches to compact round-optimal blind signatures in the CRS model which avoid including a GS proof into the final blind signature.

Another round-optimal solution, with comparable computational costs, was proposed by Seo and Cheon [316] building on work by Meiklejohn et al. [258].

Removing the CRS. Known impossibility results indicate that the design of round-optimal blind signatures in the standard model has some limitations. Lindell [244] showed that con-

currently secure (and consequently also round-optimal) blind signatures are impossible in the standard model when using simulation-based security notions. This can however be bypassed with game-based security notions, as shown by Hazay et al. [209] for non-round-optimal constructions. Fischlin and Schröder [178] showed that black-box reductions to non-interactive assumptions in the standard model are impossible for blind signature schemes under certain conditions. Roughly speaking, their result applies to schemes with three or less moves, where blindness holds statistically (or computationally if unforgeability and blindness are unrelated) and where protocol transcripts allow to verify if the user is able to derive a valid signature. Existing constructions [186, 185] bypass these results by making non-black-box use of the underlying primitives (and preventing signature-derivation checks in [186]). Unfortunately, they are impractical, as they require complexity leveraging.

At Crypto'11 Garg et al. [186] proposed the first round-optimal generic construction in the standard model. They use fully homomorphic encryption (FHE) to encrypt the message sent to the signer, who evaluates the signing circuit on the ciphertext. In addition, they use two-round witness-indistinguishable proofs (ZAPs) to remove the CRS and have the user prove that the keys have been generated honestly whereas the signer certifies honest computation. To preserve round-optimality of the scheme, the first fixed round of the ZAP is included in the signer's public key. This result can, however, only be considered as a theoretical feasibility result.

At Eurocrypt'14 Garg and Gupta [185] proposed the first efficient round-optimal blind signature constructions in the standard model. They build on Fischlin's framework using structure-preserving signatures. To remove a trusted setup, they use a two-CRS NIZK proof system (based on GS proofs), include the CRSs in the public key while forcing the signer to honestly generate the CRS. Their construction, however, requires complexity leveraging (the reduction for unforgeability needs to solve a subexponential DL instance for every signing query) and is proven secure with respect to non-uniform adversaries. Consequently, communication complexity is in the order of hundreds of KB (even at the 80-bit security level) and the computational costs (which are not even considered by the authors) seem to limit their practical application even more significantly.

Partially blind signatures. Partially blind signatures are an extension of blind signatures, which additionally allow to include a common information in a signature. Many non-round-optimal partially blind signature schemes in the ROM are based on a technique by Abe and Okamoto [19].

The latter [288] proposed an efficient construction for non-round-optimal blind as well as partially blind signatures in the standard model. Round-optimal partially blind signatures in the CRS model can again be obtained from Fischlin's framework [177]. Besides blind signatures, [63, 258, 316] also construct round-optimal partially blind signatures relying on a CRS. To date, there is—to the best of our knowledge—no round-optimal partially blind signature scheme that is secure in the standard model.

One-show anonymous credentials systems. Such systems allow a user to obtain a credential on several attributes from an issuer. The user can later selectively show attributes (or prove

relations about attributes) to a verifier without revealing any information about undisclosed attributes. No party (including the issuer) can link the issuing of a credential to any of its showings, yet different showings of the same credential are linkable. An efficient implementation of one-show anonymous credentials is Microsoft's U-Prove [79]. Baldimtsi and Lysyanskaya [40] showed that the underlying signature scheme [77] cannot be proven secure using known techniques.

To mitigate this problem, in [40] they presented a generic construction of one-show anonymous credentials in the vein of Brands' [77] approach from so-called blind signatures with attributes. They also present a scheme based on a non-round-optimal blind signature scheme by Abe [10] and prove their construction secure in the ROM.

Contributions. Besides Fischlin's generic *commit-prove* paradigm [177], there are other classes of schemes. For instance, RSA and BLS blind signatures [64, 51] follow a *randomize-derandomize* approach, which exploits the commutative or homomorphic property of the respective signature scheme. Other approaches follow the *commit-rerandomize-transform* paradigm, where a signature on a commitment to a message can be transformed into a rerandomized (unlinkable) signature on the original message [62, 190]. In contrast to these paradigms, our construction is based on a new concept, which one may call *commit-randomize-derandomize-open* approach. It does not use non-interactive proofs at all and is solely based on the recent concept of structure-preserving signature schemes on equivalence classes (SPS-EQ) [208] and commitments. We also show how to avoid a trusted setup of the commitment parameters and thus do not require a CRS.

In SPS-EQ the message space is partitioned into equivalence classes and given a signature on a message anyone can *adapt* the signature to a different representative of the same class. SPS-EQ requires that after signing a representative a signer cannot distinguish between an adapted signature for a new representative of the same class and a fresh signature on a completely random message.

In our blind-signature scheme the obtainer combines a commitment to the message with a normalization element yielding a representative of an equivalence class (*commit*). She chooses a random representative of the same class (*randomize*), on which the signer produces a signature. The obtainer then adapts it to the original representative containing the commitment (*derandomize*). This can be done without requiring the signing key and the normalization element guarantees that the obtainer can only switch back to the original representative. The blind signature is the rerandomized (unlinkable) signature for the original representative plus an opening for the commitment (*open*).

The contributions be summarized as follows:

- We propose a new approach to constructing blind signatures based on SPS-EQ in the standard model. It yields conceptually simple and compact constructions and does not rely on techniques such as complexity leveraging. Our blind signatures are practical in terms of key sizes, signature sizes, communication and computational effort (when implemented with known instantiations of SPS-EQ, a blind signature consists of 5 bilinear-group elements).

- We provide the first construction of round-optimal partially blind signatures in the standard model, which follow straightforwardly from our blind signatures and are almost as efficient.
- We generalize our blind signature scheme to message vectors, which yields one-show anonymous credentials à la "anonymous credentials light" [40]. We thus obtain one-show anonymous credentials secure in the standard model (whereas all previous ones have either no security proof or ones in the ROM).
- We also present new results in the context of SPS-EQ. We give the first construction whose security relies on a non-interactive assumption. (Unfortunately, the scheme does not have all the properties required for building blind signatures from it.) Moreover, we show how any SPS-EQ scheme can be turned into a standard SPS scheme. This transformation allows us to apply the optimality criteria by Abe et al. [15] to SPS-EQ.
- Finally, applying the impossibility result of [178], we rule out a large class of non-interactive assumptions for the unforgeability of standard-model SPS-EQ constructions.

6.8 Design Strategies for a Privacy-Friendly Austrian eID System in the Public Cloud

The Austrian eID system constitutes one major building block within the Austrian e-Government strategy, which is laid down in the Austrian e-Government Act [175] as national law. Secure authentication and unique identification of Austrian citizens – by still preserving citizens' privacy – are the main functions of the Austrian eID system. The basic building block for secure authentication and unique identification in the Austrian eID system is the Austrian citizen card [237], the official eID in Austria.

To facilitate the adoption of this eID concept at online applications, the open source module MOA-ID has been developed. Basically, MOA-ID manages the identification and authentication process based on the Austrian citizen card for service providers. Currently, the Austrian eID concept treats MOA-ID as a trusted entity, which is deployed locally 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. For instance, a centralized MOA-ID can save service providers a lot of operational and maintenance costs. However, in terms of scalability – theoretically the whole Austrian population could use this central service for identification and authentication at service providers – the existing approach is advantageous.

To bypass the issue of scalability, we propose a centralized deployment approach of MOA-ID in the public cloud. The public cloud is able to provide nearly unlimited computing resources and hence the scalability problem can easily be compensated. However, the move of a trusted service into the public cloud brings up new obstacles. In particular, MOA-ID, since now running in the public cloud, can no longer be considered a fully trusted entity. We encounter these obstacles by introducing three different approaches relying on different cryptographic mechanisms, each describing how the current Austrian eID system can be securely migrated into the public cloud. All approaches retain the workflow of the current Austrian eID system and preserve citizens'

privacy when assuming that MOA-ID acts honest but curious, i.e., operates correctly but is not trusted with respect to (data) privacy. Our first approach relies on the use of proxy re-encryption and redactable signatures, the second is based upon anonymous credentials, and the third one sets up on fully homomorphic encryption. Based on an evaluation of these three different approaches, we propose a model which can be generically used for eID identification and authentication in privacy-invasive environments such as the public cloud.

Porting the Austrian eID System to the Public Cloud. The current Austrian eID system relies on a local deployment model, where MOA-ID is deployed and operated in basically every service provider's domain. Due to that fact, MOA-ID is assumed to be trusted, i.e., it will not leak sensitive information such as the citizen's sourcePIN. The current local deployment model of MOA-ID has some benefits in terms of end-to-end security or scalability, but still some issues can be identified compared to a centralized deployment model of MOA-ID. The adoption of a centralized model may have the following advantages and disadvantages:

On the one hand, the use of one single and central instance of MOA-ID has a clear advantage for citizens as they only need to trust one specific identity provider. In addition, users could benefit from a comfortable single sign-on (SSO). On the other hand, especially service providers can save a lot of costs because they do not need to operate and maintain a separate MOA-ID installation. In addition, several different identity protocols can be supported and hence the service provider could select its favorite protocol.

Nevertheless, still some disadvantages can be identified. For instance, a single instance of MOA-ID constitutes a single point of failure or attack. Additionally, a centralized MOA-ID relies on a brokered trust model and the service provider has no direct trust relationship with the citizen anymore as it is in the local deployment model. Finally, scalability may be an issue as all citizen authentications will run through this centralized system. This is probably the main issue, as theoretically the whole Austrian population could use this service for identification and authentication at service providers. However, the issue on scalability can be tackled by moving MOA-ID into a public cloud, which is able to theoretically provide unlimited computing resources. Needless to say, a move of a trusted service into the public cloud brings up some new obstacles. For instance, assuming that MOA-ID in the cloud works correctly, it still has to be ensured that the cloud provider has no access to private citizen data during the authentication process. In general, MOA-ID in the cloud must still work equivalent to the current Austrian eID system and must still be compliant to the Austrian national law.

In order to make a migration of the Austrian eID system and MOA-ID into the public cloud possible, we have identified three approaches to adapt the existing Austrian eID system for running it in the public cloud. The adapted Austrian eID system of the respective solution will provide all functions of MOA-ID (identification, ssPIN generation, and authentication) as in the current settings, but protects citizen's privacy with respect to the cloud provider. For providing compact descriptions, we denote the SourcePIN Register Authority by SRA and the Identity Link by $\mathcal{I} = ((A_1, a_1), \ldots, (A_k, a_k))$ as a sequence of attribute labels and attribute values. Let the set of citizens be $C = \{C_1, \ldots, C_n\}$ and the set of service providers be $S = \{S_1, \ldots, S_\ell\}$ as well as the citizen's client-side middleware be denoted as M. Moreover, let us assume that

Citizen C_i wants to authenticate at service provider S_j who requires the set of attributes $\mathcal{A}_j$ from $\mathcal{I}$ and exactly one "pseudonym", i.e., the ssPIN for the sector s the service provider S_j is associated to. Additionally, recall that every citizen C_i has a signing key sk_{C_i} stored on the card and the public key pk_{C_i} is publicly available. For simplicity, we use a single security parameter κ for all schemes, although depending on the instantiation of the respective scheme different ones may be required.

Using Proxy Re-Encryption and Redactable Signatures. Here, the Identity Link $\mathcal{I}$ is modified in a way that it does no longer include the sourcePIN, but additionally all ssPINs according to all possible governmental sectors. In this augmented Identity Link $\mathcal{I}'$, every attribute a_i is encrypted using a uni-directional single-use proxy re-encryption scheme under a public key (the identity of MOA-ID) such that the corresponding private key is *not* available to MOA-ID and is only known to the SRA. Furthermore, instead of using a conventional digital signature scheme, $\mathcal{I}'$ is signed by the SRA using a redactable signature scheme such that every a_i from $\mathcal{I}'$ can be redacted. The public verification key is available to MOA-ID. Every service provider S_j obtains a key pair for the proxy re-encryption scheme when registering at the SRA. The latter entity produces a re-encryption key, which allows to re-encrypt ciphertexts intended for MOA-ID to S_j, and gives it to MOA-ID.

Using Anonymous Credentials. As above, the Identity Link $\mathcal{I}$ is augmented to $\mathcal{I}'$ in a way that it does no longer include the sourcePIN, but additionally all ssPIN's. Now, the SRA issues an anonymous credential `Cred` to every citizen for `attr` being all attributes in $\mathcal{I}'$. Essentially, a citizen then authenticates to a service provider by proving to MOA-ID the possession of a valid credential, i.e., MOA-ID checks whether the credential has been revoked or not. Note that for one-show credentials, if the entire credential `Cred` is shown to MOA-ID, this amounts to a simple lookup in a blacklist. If the credential is not revoked, MOA-ID signs the credential to confirm that it is not revoked and the citizen performs via M a non-interactive proof by revealing the necessary attributes $\mathcal{A}_j$ including the required ssPIN to S_j, who can then in turn verify the proof(s) as well as MOA-ID's signature.

Using Fully Homomorphic Encryption. This approach is a rather theoretic one and requires an FHE scheme as discussed before. The idea behind this approach is the following: The Identity Link $\mathcal{I}$ of a citizen holds the same attributes as it is in the currently deployed system (and in particular the sourcePIN), but every attribute a_i is encrypted using an FHE scheme with the above described property under MOA-ID's public key, for which MOA-ID does *not* hold the private key. Furthermore, this resulting $\mathcal{I}'$ is conventionally signed by the SRA. Then, for authentication at S_j, the resulting $\mathcal{I}'$ and the signature σ are sent to MOA-ID who checks the signature and homomorphically computes the respective ssPIN from the encrypted sourcePIN (without learning neither sourcePIN nor ssPIN). Then, for all encrypted attributes required by S_j (including the afore computed encrypted ssPIN), MOA-ID performs the "FHE re-encryption" to S_j's public key. On receiving the respective information from MOA-ID, the service provider can decrypt all attribute values.

6.9 Strengthening Authentication with Privacy-Preserving Location Verification of Mobile Phones

Smartphones are probably the most pervasive electronic devices of our time. The services they provide go far beyond communication and have become so rich and useful in daily life that most individuals can no longer think of living without them. Services increasingly make use of the geographic location of the phone to enhance the user's experience, improve usability, or strengthen the security of authentication. We propose a mechanism for using the mobile phone location as an additional authentication factor, but without all the known privacy concerns of location-based services, such as service providers creating movement and behavior patterns from the data they collect, selling this data, or losing it in a data breach. More specifically, for many security-critical scenarios, it is possible to identify a certain geographical area from which legitimate transactions are likely to occur. If the mobile phone of the user performing the transaction can be localized within this area at the time of the transaction, then this is a strong indicator that the transaction is not fraudulent. Consider, e.g., payment transactions with conventional credit cards, where the majority of credit-card fraud occurs from foreign countries geographically distant from the user's home country. To prevent this type of fraud, transactions from abroad are in many cases denied unless the user notifies the bank that she will travel, or she is contacted by the bank to validate the legitimacy of the transaction. We note that the fact that the user's phone is currently within the country where the transaction takes place can be leveraged as an authentication factor. In this scenario, it is sufficient for the credit card company to know that the phone is *somewhere* within the expected country, without learning its exact location. This idea is the basis for our proposal, as it allows for verifying the location of a user's phone without the call-center costs and without eroding the user's privacy. The bigger the geographical reference area, the better the privacy of the user can be preserved.

In addition to payment transactions, there are many other security-critical scenarios in which location can be used as an authentication factor in a privacy-preserving manner. In general, these scenarios have four entities in common: the *Service Provider* S (the payment terminal and the bank infrastructure that it connects to), the *User* U (the subscriber or cardholder), her mobile phone P, and the Mobile Network Operator M managing the mobile network to which the phone connects to.

Determining and using the location of the mobile phone does not replace authenticating the User, i.e., one aspect is ensuring that the mobile phone is at a certain location and another that the User is also there. There are existing approaches to user authentication whereby she either authenticates directly to the mobile phone or to the Service Provider, e.g., using a password or some sort of two factor authentication [291]. User authentication can alternatively take place without involving the mobile phone by requiring the user to authenticate to the physical infrastructure, e.g., with a card, a PIN, and/or a fingerprint.

Contribution We contribute a *practical, secure and privacy-preserving mechanism enabling a Service Provider to verify whether the mobile phone of a given User currently resides within a certain geographical reference area at a given time.* Our mechanism consists of having the location of the mobile phone determined by the Mobile Network Operator and certified using *anonymous*

credentials. More specifically, we contribute the design of a novel, secure and privacy-by-design mechanism fulfilling the following set of properties:

Location Unforgeability. It is infeasible for the area in which the Mobile Network Operator reports the user to be located—which we refer to as *Reported Area*—to be spoofed, tampered with, or reused later in time. In particular, the location verification mechanism should withstand strong attackers capable of fully compromising the mobile phone OS and/or its apps. In our mechanism, the potentially compromised mobile phone is not actively involved in the location determination process, since this is entirely done by the MNO, which we assume to be a trustworthy location source. Furthermore, the underlying anonymous credentials are unforgeable and resistant against tampering and replay.

User Binding. The Service Provider can rest assured that the mobile phone whose location is reported by the MNO is indeed the mobile phone of the User whom the Service Provider seeks to authenticate. In our mechanism, this property is achieved by having two credentials issued to the User and stored in her mobile phone: a Phone Credential and a User Credential. The former includes the phone identifier and is issued by the MNO after it has authenticated the mobile phone. The latter is issued by the Service Provider once it has authenticated the User and includes both the user and the phone identifiers. Binding is established when the phone generates verifiable evidence for the Service Provider attesting that the Reported Area is bound to a valid phone identifier, which is in turn bound to a certain user identifier.

Phone Identifier Privacy. Phone identifiers such as the mobile phone number, the International Mobile Equipment Identifier, or the International Mobile Subscriber Identifier are not disclosed to the Service Provider. In our mechanism, this property is achieved since information pertaining the mobile phone is only shared with the Service Provider in unintelligible form. In particular, the phone identifier included in the User Credential is carried over from the Phone Credential without revealing its value to the issuing Service Provider.

Reported Area Privacy. The Service Provider merely learns whether the mobile phone is within the area where it expects it to be—which we refer to as *Reference Area*—in a yes/no fashion, but not exactly where. If the mobile phone resides outside of the Reference Area, the Service Provider does not learn its location. In our mechanism, for achieving this property, we rely on a feature of anonymous credentials that allows proving that a given mathematical statement over certified attributes holds without disclosing their values. The mobile phone can thus prove to the Service Provider that the Reported Area is contained within the Reference Area without revealing the actual location.

Reference Area Privacy. The MNO does not learn the Reference Area. In our mechanism, this property is achieved since the Reference Area is never communicated to the MNO.

Service Unlinkability. The MNO cannot link location certification requests made by users via their mobile phones to a particular Service Provider. In our mechanism, this property is achieved because there is no direct communication between the MNO and the Service Provider, and the mobile phone never shares information concerning the Service Provider with the MNO.

Sensor Independence. The location verification process does not depend on dedicated hardware of the mobile phone. In particular, location verification does not depend on a Global

Positioning System (GPS) module or a WiFi radio being available, enabled, or active. In our mechanism, the mobile phone's lack of involvement in determining its location leads to sensor independence, which at the same time allows it to be used with any phone that is connected to a mobile network and is capable of running software to perform computations on anonymous credentials.

Usability. The User is not actively involved in the location verification process. In our mechanism this property is achieved since all steps concerning location determination and verification are designed not to require any user interaction.

6.10 Blind Attribute-Based Encryption and Oblivious Transfer with Fine-Grained Access Control

Oblivious transfer [301] (OT) is a cryptographic primitive that allows the construction of privacy-preserving databases, i.e., databases that allow users to query and retrieve information anonymously and without the data holder learning which information is accessed. Thanks to these properties, oblivious transfer has been proposed to solve privacy-related problems in a wide variety of applications that involve sensitive information, such as medical databases, patent searches or location-based services [273, 232].

However, current regulations (e.g., HIPAA) require data holders to enact strict accounting procedures in order to ensure that data is only accessed by authorized parties. Therefore, it is desirable to provide databases based on OT with privacy-preserving access control mechanisms that allow the data holder to enforce access control policies. The current trend towards distributing and outsourcing sensitive information (e.g. Microsoft HealthVault) makes this requirement even more compelling.

Privacy-preserving access control mechanisms for OT must define, for each record of data, an access control policy that determines the attributes, rights or roles that a user must possess in order to be granted access. However, this should be done without harming the privacy properties of OT, i.e., users must remain anonymous towards the data holder and must still be able to hide from him both their attributes and the records they access. Additionally, the access control mechanism should be flexible enough to allow the data holder to apply a wide variety of access control policies.

Previous work K-out-of-N [122] OT is a two-party protocol between a sender and a receiver. The sender receives as input N messages $(m_1, \ldots, m_N)$, while the receiver gets K selection values $(\sigma_1, \ldots, \sigma_K)$. As output, the receiver gets the messages $(m_{\sigma_1}, \ldots, m_{\sigma_K})$. Sender privacy requires that the receiver gets no information on the other messages, while receiver privacy requires that the sender does not learn any information on $(\sigma_1, \ldots, \sigma_K)$. In adaptive OT, the receiver chooses σ_i $(i \in [1, K])$ after outputting message $m_{\sigma_{i-1}}$, while, in non-adaptive OT, the selection values $(\sigma_1, \ldots, \sigma_K)$ are chosen before outputting any message.

Oblivious transfer with access control (OTAC) [156, 94] allows the sender to control access to the messages. The sender receives as input $(m_1, \mathbb{P}_1, \ldots, m_N, \mathbb{P}_N)$, where $(\mathbb{P}_1, \ldots, \mathbb{P}_N)$ are access

control policies for each of the messages. Each receiver possesses a set of attributes $\mathbb{A}$, which is certified by a credential issuer, and is able to obtain the message m_i only if $\mathbb{A}$ satisfies $\mathbb{P}_i$. Receiver privacy requires that the sender does not get any information on $\mathbb{A}$.

Coull et al. [156] propose a scheme in which access control policies are described by a state graph, such that each state allows access to a subset of the records. In the initialization phase, the database records are encrypted on input an index from 1 to N. Additionally, each user obtains a credential that binds her to a state in the graph. To access a record, the user proves possession of her credential and proves that the index of the record is included in the subset of indices defined by her state. Each access to the database determines a transition from one state to another, which depends on the record accessed. The user obtains a new credential that binds her to her new state. Access control is enforced by limiting the possible transitions between states.

This scheme has the nice properties that it can be applied to different OT schemes and that it permits changing the policies without needing to rerun the initialization phase of the OT scheme. However, as noted by Camenisch et al. [94], the scheme is inefficient for two reasons. First, users must obtain a new credential each time they access the database. Second, a large class of access control policies cannot be efficiently expressed via state graphs.

Camenisch et al. [94] modify the OT scheme in [122] to provide access control. First, each user obtains a credential that certifies that she possesses some attributes. In the initialization phase, the sender encrypts each record of data on input a set of attributes. To access a record, the user must prove that she possesses a credential that contains all the record's attributes. The access control policy class that can be expressed is therefore limited to conjunction of attributes. To allow expressing disjunction, the database holder can replicate the record, increasing thus the database size. For example, let (a_1, a_2, a_3, a_4) be attributes and $(a_1 \wedge a_2) \vee (a_3 \wedge a_4)$ be the policy for record m. The sender can input to the database two records $m_1 = m_2 = m$, where the policy for m_1 is $(a_1 \wedge a_2)$ and, for m_2, $(a_3 \wedge a_4)$. A more recent work [12], also limited to conjunctive policies, achieves universal composability.

After the work by Camenisch et al. [94], other schemes aim at achieving better efficiency or at supporting more expressive access control policies. To achieve better efficiency, the access control policies in the protocol in [344] are limited to proving possession of one attribute. In the protocol in [205], the user authenticates herself using an anonymous credential, but the messages are not associated to access control policies.

To support more expressive access control policies, the protocol in [342] employs the fuzzy identity based encryption scheme by Sahai and Waters [311] in order to support threshold policies. In threshold policies, each record is associated with a set of attributes and a user is granted access if the number of attributes in the policy that she possesses is over a threshold.

Other protocols go further and employ ciphertext policy attribute based encryption (CPABE) [307, 348, 204, 343, 199, 200]. Generally speaking, these protocols work as follows. In the initialization phase, each record of data is associated with an access control policy. The sender encrypts the records of data under their respective policies via the CPABE scheme. To access a record, the user must hold a secret key whose attributes fulfill the record's access control policy.

In some schemes [204, 200, 348], users must reveal their attributes in order to obtain a secret key of the CPABE scheme. We note that in [94] the credential issuer does not necessarily learn the users' attributes. Namely, the credential issuer can certify the attributes by letting the user prove in zero-knowledge statements about them.

To solve this problem, an earlier version of this work [307] and other schemes [343, 199] propose the use of a blind key extraction protocol for the CPABE scheme. The protocol in [199] employs blind CPABE but the protocol involves only a sender and a receiver, who obtains the secret key for the CPABE scheme engaging with a sender in a blind key extraction protocol. However, the attributes of the user are not certified, so it is unclear how access control is actually enforced.

The protocol in [343] employs blind CPABE along with a credential scheme. The concrete blind CPABE proposed supports policies that are described by conjunctions and disjunctions of attributes, but threshold policies are not efficiently supported. Additionally, the parameters of the CPABE scheme grow with the number of possible attributes and the CPABE is not proven to be committing, a property that ensures that a corrupt sender cannot compute malformed ciphertexts that decrypt to different messages. Finally, although a concrete blind ABE is proposed, no generic way of constructing blind CPABE with access control protocols for any CPABE scheme is proposed.

Prior to [343, 199], an earlier version of this work [307] proposed a construction of OT with access control from any CPABE scheme, oblivious transfer and anonymous credentials scheme. The concrete CPABE scheme used supports efficiently any policy that can be expressed by a monotonic access structure, which includes conjunctive, disjunctive and threshold policies, and, although the ciphertext size grows with the size of the access structure, the size of the public parameters is independent of the number of attributes. In [307], a generic way of building blind CPABE with access control from CPABE is proposed, which employs oblivious transfer and anonymous credentials. However, although security for OT and anonymous credentials is defined via an ideal functionality, the proposed protocol does not employ these functionalities as building blocks in a hybrid protocol, which hinders proving security of the construction.

In a related line of work, Camenisch et al. [95] propose a scheme where access control policies are hidden from the user. However, the access control policy class is limited to conjunction of attributes. In [91], Camenisch et al. improve the work in [95]. They propose an oblivious transfer with hidden access control policies scheme based on the hidden ciphertext-policy attribute based-encryption by Nishide et al. [281]. The access control policy class supported is described as follows. Let n be the number of attribute types. The set S_i $(i \in [1, n])$ contains the possible values that an attribute s_i of type i can have. A receiver possesses a set of attributes $(s_1, \ldots, s_n)$ such that $s_i \in S_i$ $(i \in [1, n])$. An access control policy is given by $(W_1, \ldots, W_n)$, where $W_i \subseteq S_i$ is a subset of S_i. The policy is satisfied if, for $i = 1$ to N, $s_i \in W_i$. As can be seen, to express disjunction of attributes of different categories, the sender must also replicate messages. Moreover, the credential issuer necessarily learns the attributes to be able to compute a secret key of the CPABE scheme.

Our Contribution We propose a non-restricted and a restricted OTAC scheme. In a non-restricted OTAC scheme, the receiver can obtain in one transfer phase all the messages whose

access control policy is fulfilled by the receiver's attributes. In a restricted OTAC scheme, the receiver can only obtain one message per transfer phase.

Our schemes employ CPABE, a non-adaptive OT functionality, and a simple anonymous credentials functionality referred to as anonymous attribute authentication. The restricted OTAC construction also employs an adaptive OT functionality.

The sender employs committing CPABE to encrypt the messages under their respective access control policies. We define the committing property, which ensures that malformed ciphertexts cannot yield different messages.

To access the messages, a receiver gets a CPABE secret key for her attributes from the sender by using a blind key extraction with access control protocol. We provide a generic construction of a blind key extraction with access control protocol that works for any committing CPABE scheme that fulfills an additional property we call key separation. Basically, key separation guarantees that a secret key for attributes $(a_1, \ldots, a_L)$ consists of $(sk_{a_1}, \ldots, sk_{a_L}, sk'_{\mathbb{A}})$, where each sk_{a_l} is computed on input the attribute a_l only, and $sk'_{\mathbb{A}}$ is computed without knowledge of any attribute.

Our generic construction of blind key extraction with access control works in a hybrid model that employs two novel ideal functionalities for non-adaptive OT and for anonymous attribute authentication. We consider a universe of attributes $[1, L_{uni}]$, where L_{uni} is the size of the universe. The sender inputs a secret key $(sk_{a_1}, \ldots, sk_{a_{L_{uni}}})$ to the non-adaptive OT functionality, where each sk_{a_l} acts as a message to be transferred. The issuer issues some attributes $(a_1, \ldots, a_L)$ to a receiver through the anonymous attribute authentication functionality. In order to obtain a secret key for the attributes issued, the receiver must prove to the sender that these attributes were issued by the issuer through the anonymous attribute authentication functionality. Additionally, the receiver must prove that the selection values $(\sigma_1, \ldots, \sigma_K)$ sent as input to the non-adaptive OT functionality equal the attributes sent as input to the anonymous attribute authentication functionality.

However, existing ideal functionalities for non-adaptive OT do not allow for that, i.e., in a hybrid protocol where the non-adaptive OT functionality is used along with other functionalities, existing non-adaptive OT functionalities do not provide a means that allows the sender to check whether the selection values input to the non-adaptive OT functionality equal those input to other functionalities. To solve this problem, we propose a novel non-adaptive OT functionality and a novel anonymous attribute authentication functionality where the receiver sends committed inputs. The sender receives those commitments from the functionalities and checks their equality. Under the binding property of the commitment scheme, it holds that the input to both functionalities is equal. The idea of using a commitment scheme for this purpose is taken from our revocation construction in Section 7.1.

We propose a concrete committing CPABE with key separation scheme based on the CPABE scheme in [57]. We also show constructions that realize our novel non-adaptive OT and anonymous attribute authentication functionalities. We define security of non-restricted and restricted OTAC protocols in the universal composability framework and we prove that our protocols for non-restricted and restricted OTAC realize their respective functionalities. The constructions

we propose do not achieve universal composability because, for the sake of efficiency, we employ the Fiat-Shamir transform [176] to instantiate non-interactive zero-knowledge proofs.

6.11 Privacy-Preserving Smart Metering Revisited

In privacy-preserving billing a meter measures a user's consumption of some utility or service and service providers apply fine-grained tariff policies, i.e., policies that require detailed and frequent consumption measurements, in order to determine the bill.

A classical example is smart metering of electricity, water and gas [256]. In this setting, utility providers install smart meters in households in order to measure user consumption. Smart meters provide meter readings to the service provider. These readings are used by the service provider to calculate the bill under the tariff policy. The tariff policy may be complex, e.g., by applying a different rate depending on the time of consumption or on whether the consumption measurement reaches a threshold.

Other examples are electronic toll collection [211] and pay-as-you-drive car insurance [72]. In these cases, drivers install a meter in their cars that reports to the service provider which roads are used and when. Typical tariff policies apply different rates depending on the type of road (e.g. motorway, street), the time of the day (e.g. day or night), or even the speed of the vehicle.

In all the settings above, billing poses a threat to user privacy. Meters report fine-grained readings to the service provider, which potentially discloses sensitive information. For example, electricity smart-meter readings reveal when users are at home and the electrical appliances they use [28], and electronic toll collection and "pay as you drive" insurance reveal the driver's whereabouts [286, 38, 330].

In privacy-preserving billing protocols, meters do not send consumption measurements to the service provider. Instead, the computation of the bill is done locally and only the amount to be paid is revealed to the service provider.

Privacy-preserving billing protocols, in particular those which employ meters that are not tamper-resistant, involve mechanisms to ensure that users report meter readings correctly, such as random spot-checks in the electronic toll collection protocol in [38, 257, 299].

The protocols that use tamper-resistant meters either perform the bill calculation in the meter or outsource it to an untrusted platform to keep the meters simple. In [330], the bill calculation is performed inside the tamper-resistant meter. In contrast, in [305] the tamper-resistant meter outputs signed meter readings to a user application. At the end of a billing period, the user application employs the tariff policy sent by the service provider and the signed readings obtained from the meter to calculate the bill. The user application reveals to the service provider only the total bill, along with a proof that the computation of the bill is correct. This proof does not reveal any additional information on the meter readings. The approach in [305] has the advantage that it allows to minimize the trusted computing base and that it avoids the need to update tamper-resistant meters when the tariff policy changes. Practical implementations of these protocols have been shown in [180].

Our contribution. We revisit the work of [305] and improve it in two ways. First, we generalize the security model in [305] to consider multiple meters and multiple users. Second, we propose a privacy-preserving billing protocol for our model that, in comparison to the protocol in [305], improves efficiency for policies described by splines.

The security model in [305] considered a setting where a meter communicates only with one user, and a user communicates only with one meter, i.e., there is a one-to-one relation between users and meters. This is insufficient for some smart metering applications. For example, consider an apartment building with central heating. Each apartment is provided with a smart meter that measures the electricity consumption of its tenants. Additionally, another meter measures the electricity consumption of each of the tenants with respect to the central heating system. To model this setting adequately, it is necessary to both allow a meter to send meter readings to multiple users, and to allow a tenant to receive meter readings from more than one meter.

Therefore, we propose an ideal functionality $\mathcal{F}_{\text{BIL}}$ for privacy preserving billing protocols that considers multiple meters and multiple users. In addition to that, $\mathcal{F}_{\text{BIL}}$ has the following main differences in comparison to the functionality for smart metering described in [305].

- $\mathcal{F}_{\text{BIL}}$ includes an interface through which the service provider sends a list of meters to a user at each billing period. The meter readings received from the meters in the list must be employed by the user to perform the bill calculation for that billing period.
- $\mathcal{F}_{\text{BIL}}$ includes an interface that allows meters to signal the end of a billing period and to report to the users the number of meter readings that were sent during the billing period. This necessary interface was omitted in the functionality in [305].
- $\mathcal{F}_{\text{BIL}}$ models explicitly the communication with the simulator $\mathcal{S}$. $\mathcal{S}$ needs this communication in order to provide a simulation for the adversary in the security proof.
- $\mathcal{F}_{\text{BIL}}$ allows any verifying party (and not just the service provider) to verify the bill to be paid. This may be useful in case of dispute between the meter and the service provider.
- $\mathcal{F}_{\text{BIL}}$ models the cases in which corrupt users collude with corrupt meters and/or with the service provider.

We propose a privacy-preserving billing protocol that realizes our functionality $\mathcal{F}_{\text{BIL}}$ and thus allows a meter to send meter readings to multiple users, and users to employ meter readings from multiple meters in the computation of a bill. In a nutshell, our protocol work as follows. At each billing period, the provider registers a signed tariff policy. Tariff policies are of the form $Y : (c, t) \to p$, where c is the consumption measurement, t is the time of consumption, and p is the price. The provider also sends to each user a signed list of meters. Meters send signed meter readings to users and a signed "end of billing period" message that contains the number of meter readings sent from the meter to the user at that billing period. The user application calculates the bill and computes a zero-knowledge proof of knowledge of its correctness. This zero-knowledge proof involves proofs of signature possession that demonstrate that the correct tariff policy is used to compute the price for each of the signed meter readings.

In [305], it is shown how to sign different types of tariff policies: a linear policy that multiplies each reading by a price per unit of consumption and a cumulative policy that divides the consumption range in intervals and applies a different price per unit to each interval. Additionally, it is mentioned that, in general, a tariff policy may be described by a polynomial for each interval. (Other functions can be approximated by polynomial splines.) Although the protocol in [305] provides efficient zero-knowledge proofs for the linear and cumulative policies, the cost of a zero-knowledge proof of a tariff policy described by a polynomial grows with the polynomial degree.

Our privacy-preserving billing protocol employs the same technique in [305] to sign linear and cumulative policies, and employs a new method for tariff policies described by splines. Consider the following tariff policy as example. A day is divided into L time intervals. For each time interval, the price to be paid for the consumption c is given by a spline:

$$Y(c,t) = \left\{ \begin{array}{ll} \Phi_1(c) & \text{if } t \in [t_1, t_2) \\ \vdots & \vdots \\ \Phi_L(c) & \text{if } t \in [t_L, t_{L+1}) \end{array} \right\}$$

Each spline $\Phi_l(c)$ $(l \in [1, L])$ is defined as follows.

$$\Phi_l(c) = \left\{ \begin{array}{ll} \phi_1(c) & \text{if } c \in [c_1, c_2) \\ \vdots & \vdots \\ \phi_M(c) & \text{if } c \in [c_M, c_{M+1}) \end{array} \right\}$$

Therefore, for a meter reading (c, t), the price to be paid is defined by the polynomial $\phi_m(c)$ such that $c \in [c_m, c_{m+1})$ that belongs to the spline $\Phi_l(c)$ associated to the time interval $[t_l, t_{l+1})$ such that $t \in [t_l, t_{l+1})$.

Alternatively, one can consider consumption bands, i.e. if a user's consumption is below a certain threshold he may get a better price at peak hours. For each consumption band, the price to be paid at a certain time of day t is given by a spline where the polynomials take the time as input.

Our method to sign a tariff policy given by splines employs the polynomial commitment scheme of [228]. In a nutshell, the service provider computes polynomial commitments C to each of the polynomials in the tariff policy for the billing period bp. Additionally, the service provider computes, for each polynomial commitment, a signature on $[bp, C, t_{l-1}, t_l, c_{m-1}, c_m]$. The service provider sends the polynomial commitments and the signatures to the users, together with the polynomials. To prove in zero knowledge that the price calculated for a meter reading is correct, the user evaluates the polynomial on input the consumption to compute the price, and then proves possession of a witness for the polynomial commitment that shows that the price is the correct evaluation of the committed polynomial. The size of this proof of witness possession is independent of the polynomial degree. Additionally, the user proves possession of a signature on the polynomial commitment, and proves that the values of consumption and time in the meter reading lie within the respective intervals in the signature.

Our use of polynomial commitments is somewhat different from their common use. In our scheme, the service provider computes polynomial commitments and sends them to the user

together with the polynomials. Therefore, we do not need the hiding property of commitments. However, we need the binding property because the polynomial commitments are employed by the user to prove in zero-knowledge that prices are computed following the polynomials that define the tariff policy.

We show that our protocol realizes $\mathcal{F}_{\text{BIL}}$. Unlike [305], we analyze all the possible collusion scenarios. Additionally, we consider the case in which the provider and the meters are corrupt but do not have a side communication channel between them. We show that, in this case, our protocol is collusion-free in the sense of [238] and prevents corrupt meters from disclosing the meter readings to the provider or another corrupt verifying party.

We discuss how our protocol compares to other possible approaches for the design of privacy-preserving billing protocols. Concretely, we discuss the use of regulations and codes of conduct, trusted parties, techniques to reduce variability, data anonymization methods, differential privacy, verifiable computing, and secure two-party and multi-party computation.

We note that our protocol is not only useful for billing, but, in general, allows to prove correctness of any computation on meter readings. This is important in settings such as smart metering, where meter readings are not only used for the sake of billing but also for consumption forecasting or profiling. For these other purposes, protocols that support complex computations on meter readings are necessary.

6.12 Computing on Authenticated Data

In tandem with recent progress on computing *any function* on encrypted data, e.g., [189, 334, 322], this work explores computing on unencrypted signed data. In the past few years, several independent lines of research touched on this area:

- Quoting/redacting: [324, 224, 32, 262, 203, 84, 82, 83] Given Alice's signature on some message m anyone should be able to derive Alice's signature on a subset of m. Quoting typically applies to signed text messages where one wants to derive Alice's signature on a substring of m. Quoting can also apply to signed images where one wants to derive a signature on a subregion of the image (say, a face or an object) and to data structures where one wants to derive a signature of a subset of the data structure such as a sub-tree of a tree.

- Arithmetic: [234, 349, 134, 70, 188, 69, 68, 340] Given Alice's signature on vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k \in \mathbb{F}_p^n$ anyone should be able to derive Alice's signature on a vector $\mathbf{v}$ in the linear span of $\mathbf{v}_1, \ldots, \mathbf{v}_k$. Arithmetic on signed data is motivated by applications to secure network coding [181]. We show that these schemes can be used to compute authenticated linear operations such as computing an authenticated weighted sum of signed data and an authenticated Fourier transform. As a practical consequence of this, we show that an untrusted database storing signed data (e.g., employee salaries) can publish an authenticated average of the data without leaking any other information about the stored data. Recent constructions go beyond linear operations and support low degree polynomial computations [68].

- Transitivity: [308, 260, 52, 213, 53, 318, 345, 277] Given Alice's signature on edges in a graph G anyone should be able to derive Alice's signature on a pair of vertices (u, v) if and only if there is a path in G from u to v. The derived signature on the pair (u, v) must be indistinguishable from a fresh signature on (u, v) had Alice generated one herself [260]. This requirement ensures that the derived signature on (u, v) reveals no information about the path from u to v used to derive the signature.

We put forth a general framework for computing on authenticated data that encompasses these lines of research and much more. While prior definitions mostly contained artifacts specific to the type of malleability they supported and, thus, were hard to compare to one another, we generalize and strengthen these disparate notions into a single definition. This definition can be instantiated with any predicate, and we allow repeated computation on the signatures (e.g., it is possible to quote from a quoted signature.) During our study, we realized that the "privacy" notions offered by many existing definitions are, in our view, insufficient for some practical applications. We therefore require a stronger (and seemingly a significantly more challenging to achieve) property called *context hiding*. Under this definition, we provide two generic solutions for computing signatures on any univariate, closed predicate; however, these generic constructions are not efficient. We also present efficient constructions for three problems: quoting substrings, a subset predicate, and a weighted average over data (which captures weighted sums and Fourier transforms). Our quoting substring construction is novel and significantly more efficient than the generic solutions. For the problems of subsets and weighted averages, we show somewhat surprising connections to respective existing solutions in attribute-based encryption and network coding signatures.

A general framework. Let $\mathcal{M}$ be some message space and let $2^{\mathcal{M}}$ be its powerset. Consider a predicate $P : 2^{\mathcal{M}} \times \mathcal{M} \to \{0, 1\}$ mapping a set of messages and a message to a bit. Loosely speaking we say that a signature scheme supports computations with respect to P if the following holds:

> Let $M \subset \mathcal{M}$ be a set of messages and let m' be a *derived* message, namely m' satisfies $P(M, m') = 1$. Then there exists an efficient procedure that can derive Alice's signature on m' from Alice's independent signatures on all of the messages in M.

For the quoting application, the predicate P is defined as $P(M, m') = 1$ iff m' is a quote from the set of messages M. Here we focus on quoting from a single message m so that P is false whenever M contains more than one component, and thus use the notation $P(m, m')$ as shorthand for $P(\{m\}, m')$. The predicate P for arithmetic computations essentially says that $P\big((\mathbf{v}_1, \ldots, \mathbf{v}_k), \mathbf{v}\big)$ is true whenever $\mathbf{v}$ is in the span of $\mathbf{v}_1, \ldots, \mathbf{v}_k$.

We emphasize that signature derivation can be iterative. For example, given a message-signature pair (m, σ) from Alice, Bob can publish a derived message-signature pair (m', σ') for an m' where $P(m, m')$ holds. Charlie, using (m', σ'), may further derive a signature σ'' on m''. In the quoting application, Charlie is quoting from a quote, which is perfectly fine.

Security. We give a clean security definition that captures two properties: unforgeability and context hiding. We briefly discuss each of them.

- Unforgeability captures the idea that an attacker may be given various derived signatures (perhaps iteratively derived) on messages of his choice. The attacker should be unable to produce a signature on a message that is not derivable from the set of signed messages in his possession. E.g., suppose Alice generates (m, σ) and gives it to Bob who then publishes a derived signature (m', σ'). Then an attacker given (m', σ') should be unable to produce a signature on m or on any other message m'' such that $P(m', m'') = 0$.
- Context hiding captures an important privacy property: a signature should reveal nothing more than the message being signed. In particular, if a signature on m' was derived from a signature on m, an attacker should not learn anything about m other than what can be inferred from m'. This should be true even if the original signature on m is revealed. For example, a signed quote should not reveal anything about the message from which it was quoted, including its length, the position of the quote, whether its parent document is the same as another quote, whether it was derived from a given signed message or generated freshly, etc.

We note that notions such as group or ring signatures [144, 50, 117, 65, 309] have considered the problem of hiding the identity of a signer among a set of users. Context hiding ensures privacy for the data rather than the signer. Our goal is to hide the legacy of how a signature was created.

Efficiency. We require that the size of a signature, whether fresh or derived, depend only on the size of the object being signed. This rules out solutions where the signature grows with each derivation.

Generic Approaches. We begin with two generic constructions that can be inefficient. They apply to *closed, univariate* predicates, namely predicates $P(M, m')$ where M contains a single message (P is false when $|M| > 1$) and where if $P(a, b) = P(b, c) = 1$ then $P(a, c) = 1$. The first construction uses any standard signature scheme S where the signing algorithm is deterministic. (One can enforce determinism using PRFs [192].) To sign a message $m \in \mathcal{M}$, one uses S to sign each message m' such that $P(m, m') = 1$. The signature consists of all these signature components. To verify a signature for m, one checks the signature component corresponding to the message m. To derive a signature m' from m, one copies the signature components for all m'' such that $P(m', m'') = 1$. Soundness of the construction follows from the security of the underlying standard scheme S and context hiding from the fact that signing in S is deterministic.

Unfortunately, these signatures may become large consisting up to $|\mathcal{M}|$ signature components — impacting both the signing time and signature size. Our second generic construction alleviates the space burden by using an RSA accumulator. The construction works in a similar brute force fashion where a signature on m is an accumulator value on all m' such that $P(m, m') = 1$. While this produces short signatures, the time component of both verification and derivation are even

worse than the first generic approach. Thus, these generic approaches are too expensive for most interesting predicates.

Our Quoting Construction. We turn to more efficient constructions. First, we set out to construct a signature for quoting *substrings*[3], which although conceptually simple is non-trivial to realize securely. As an efficiency baseline, we note that the brute force generic construction of the quoting predicate would result in n^2 components for a signature on n characters. So any interesting construction must perform more efficiently than this. We prove our construction selectively secure.[4] In addition, we give some potential future directions for achieving adaptive security and removing the use of random oracles.

Our construction uses bilinear groups to link different signature components together securely, but in such a way that the context can be hidden by a re-randomizing step in the derivation algorithm. A signature in our system on a message of length n consists of $n \lg n$ group elements; intuitively organized as $\lg n$ group elements assigned to each character. To derive a new signature on a substring of ℓ characters, one roughly removes the group elements not associated with the new substring and then re-randomizes the remaining part of the signature. This results in a new signature of $\ell \lg \ell$ group elements. The technical challenge consists in simultaneously allowing re-randomization and preserving the "linking" between successive characters. In addition, there is a second option in our derive algorithm that allows for the derivation of a short signature of $\lg \ell$ group elements; however the derive procedure cannot be applied again to this short signature. *Thus, we support quoting from quotes, and also provide a compression option which produces a very short quote, but the price for this is that it cannot be quoted from further.*

Computing Signatures on Subsets and Weighted Averages. Our final two contributions are schemes for deriving signatures on subsets and weighted averages on signatures. Rather than create entirely new systems, we show connections to existing Attribute-Based Encryption schemes and Network Coding Signatures. Briefly, our subset construction extends the concept of Naor [67] who observed that every IBE scheme can be transformed into a standard signature scheme by applying the IBE KeyGen algorithm as a signing algorithm. Here we show an analog for known Ciphertext-Policy (CP) ABE schemes. The KeyGen algorithm which generates a key for a set S of attributes can be used as a signing algorithm for the set S. For known CP-ABE systems [57, 239, 338] it is straightforward to derive a key for a subset S' of S and to re-randomize the signature/key. To verify a signature on S we can apply Naor's signature-from-IBE idea and encrypt a random message X to a policy that is an AND of all the attributes in S and see if the signature can be used as an ABE key to decrypt to X.

[3] A substring of $x_1 \ldots x_n$ is some $x_i \ldots x_j$ where $i, j \in [1, n]$ and $i \leq j$. We emphasize that we are not considering *subsequences*. Thus, it is *not* possible, in this setting, to extract a signature on "I like fish" from one on "I do not like fish".

[4] Following an analog of [130], selective security for signatures requires the attacker to give the forgery message before seeing the verification key.

6.13 (Un)linkable Pseudonyms for Governmental Databases

When data is collected and maintained on a large scale, that data often does not reside in a single database but is distributed over several databases and organisations, each being responsible for a particular aspect in the overall system. To still allow for collaborative operations and data exchange among the different entities, related data is then indexed with an identifier that is unique in the entire system.

An important example of such a distributed setting is a state-controlled social security system maintaining various sets of personal data. Therein users can interact with different entities, such as different health care providers, public and private pension funds, or tax authorities. All these entities can act autonomously in their respective domains and keep individual records for the users they interact with. In certain scenarios, the different entities also have to exchange data about particular users. For instance, assume a health care provider offers special discounts for people with low income and tax authorities store information about users' salaries. Then, to verify whether a user is eligible for a discount, the health care system together with the tax authority should be able to check if the user satisfies the criteria.

Global Identifiers Currently, the probably most prominent approach to enable such decentralized data management is to use unique global identifiers among all entities. In the context of social security systems, this is for instance implemented in the US, Sweden, and Belgium. Here, each citizen gets assigned a unique and nation-wide social security number. The advantage of this approach is that it naturally allows all entities within the system to correlate their individually maintained records. However, having such a unique identifier for each user is also a significant privacy threat: When data is lost or stolen, also any adversary obtaining the data can use the unique identifier to link all the different data sets together. Also, interactions of users with different entities become easily traceable.

Thus, the impact of security breaches is rather severe, which in turn, makes the data maintained by the individual entities a lucrative target for data thieves. In addition, as all entities can trivially link their records, the data exchange can hardly be controlled and authorized. However, in particular in the case of a social security system, a certain control to supervise and, if necessary, limit the data flow is usually desired. For instance in the Belgium system currently a central authority called "crossroads bank for social security" (CBSS) [2], serves as hub for all data exchange. Whenever social and private entities want to exchange data based on the global identification number, they have to request explicit authorization from the CBSS, as enforced by national law. From a privacy point of view, though, this added controllability makes the system even worse, as now a central authority learns which requests are made for which user. In a social security system, those requests can reveal quite sensitive information itself. For instance, in the example outlined above, the central authority would learn from the requests which persons suffer from health issues and probably have low or no income, even if it has no access to the health and tax records itself. Also in terms of security it still assumes that all entities behave honestly and do not correlate their records without approval of the central authority.

Pseudonyms & Controlled Conversion Having such a central authority actually allows for a more privacy-friendly solution. Namely, a central authority (that we call *converter*) could derive and distribute entity-specific identifiers (aka *pseudonyms*), in a way that pseudonyms known by different entities can only be linked with the help of the converter. Thus, it would then even be technically enforced that different entities have to request permission from the converter, as without its help they would not be able to connect their records anymore. Of course, the latter argument only holds if the data sets maintained by the entities do not contain other unique identifying information which allows linkage without using the pseudonyms.

Such a pseudonymous identification system clearly improves the controllability of the data exchanges and also avoids imposing a unique identifier that makes the user traceable by default. Both are significant advantages compared with the solution where only a single global identifier is used throughout the entire system. However, as now the converter is indeed required in every request it yields a powerful entity that still must be trusted to not exploit the information it gathers.

Existing Solutions In existing solutions, the need to fully trust the converter seems in fact inherent. A similar pseudonymous framework using a central converter is for instance described by Galindo and Verheul [184]. Therein, the converter computes a pseudonym $nym_{i,A}$ based on main identifier uid_i and for server $\mathcal{S}_A$, as $nym_{i,A} = \mathsf{Enc}(k_A, uid_i)$, where Enc is a blockcipher and k_A a symmetric key that the converter has chosen for $\mathcal{S}_A$, but is only known to the converter. When an entity $\mathcal{S}_A$ then wishes to request some data for $nym_{i,A}$ from another entity $\mathcal{S}_B$, it sends the pseudonym to the converter. The converter then decrypts $nym_{i,A}$ to obtain uid_i and derives the local pseudonym $nym_{i,B}$ by computing $nym_{i,B} = \mathsf{Enc}(k_B, uid_i)$ for the key k_B it had chosen for $\mathcal{S}_B$. Thus, here the converter is necessarily modeled as a trusted third party, as it always learns the generated pseudonyms, the underlying uid_i and also has full control over the translations it provides (i.e., a corrupt converter could transform pseudonyms arbitrarily).

Another example is the Austrian eID system [1], which is one of the few eID solutions that allows one to derive entity-specific pseudonyms from the unique social security number. However, it currently only supports that unlinkable pseudonyms are created by the users themselves, but it does not consider a central authority that can provide a conversion service on a large scale. It is easy to imagine though, how such a converter could be realized. Roughly, a pseudonym $nym_{i,A}$ is computed as $\mathsf{H}(\mathsf{Enc}(k, uid_i)||\mathcal{S}_A)$, i.e., the encrypted main user identifier uid_i and the identifier of the respective entity $\mathcal{S}_A$ are concatenated and the hash value of both yields the pseudonym. Here, the key k is a global key that is used for all pseudonyms, but is again only known to the converter. In order to enable conversions between pseudonyms, the converter could simply keep a table with the related hash values and then perform the conversion based on looking up the corresponding value.

Hereby, the trust requirements for the converter can actually be reduced if one considers generation and conversion of pseudonyms as two different tasks. Then, only the entity responsible for pseudonym generation would have to know the key k under which the user identifiers are encrypted, whereas the converter merely keeps the hash table with the related pseudonyms. The converter would then only know which pseudonyms belong together, but can not determine for

which particular user they are standing for. Thus, also during conversion, a malicious converter does not learn the particular user for which a conversion is requested anymore, but only his pseudonym.

However, the converter can still link all requests that are made for the same (unknown) user. As each query usually leaks some context information itself, being able to link all that information together might still allow the converter to fully identify the concrete user behind a pseudonym. For instance, regular queries for the same pseudonym to the pension fund might indicate that the person behind the pseudonym is older than 60 years, and queries to entities that are associated with a certain region such as local municipalities further reveal the place that person might live in.

Via the comparable CBSS authority in Belgium, several hundreds of million messages are exchanged every year, with a peak of 806 million messages in 2009. Using those values as a reference for the social security use case, one has to assume that a converter learning "only" the requests and their relation would still obtain a significant amount of context data. How context information and meta data can be leveraged to fully de-anonymize pseudonymized data sets, was recently impressively demonstrated for "anonymized" credit card transactions [161] and in the Netflix and AOL incidents [274, 42].

Thus, from a privacy and a security perspective it is clearly desirable to minimize the information a converter can collect as much as possible. This means, the converter should not even learn which requests relate to which pseudonyms.

Other Related Work There exists a line of work on reversible pseudonymization of data records, in particular in the eHealth context, aiming at de-sensitizing patient records [7, 276, 160, 170, 298]. The main focus in these works is to derive pseudonyms from unique patient identifiers, such that the pseudonyms do not reveal any information about the patient anymore, yet allow de-anonymization by a trusted party (or a combination of several semi-trusted parties). However, in all solutions, pseudonym generation must be repetitively unambiguous to preserve the correlation between all pseudonymized records. Consequently, data exchange is trivial and does not require a converter. Thus, pseudonyms are linkable by default, whereas our approach is the opposite: pseudonyms should be *unlinkable by default*, yet preserve the correlation which allows to re-establish the linkage only if necessary via a (potentially untrusted) converter.

Our Contribution We tackle the challenge of enabling privacy-friendly yet controlled data exchange in a decentralized system. That is, we propose an (un)linkable pseudonym system where a converter serves as central hub to ensure controllability. The converter establishes individual pseudonyms for each server derived from a unique main identifier that every user has, but without learning the derived pseudonyms. The converter is still the only authority that can link different pseudonyms together, but it does not learn the particular user or pseudonym for which such a translation is requested. The converter cannot even tell if two data exchanges were done for the same pseudonym or for two different ones. Thus, the only information the converter still learns is that a server $\mathcal{S}_A$ wants to access data from a server $\mathcal{S}_B$. We consider this to be the right amount of information to balance control and privacy. For instance for the use case of a

social security system, it might be allowed that the health care provider can request data from the tax authority but should not be able to access the criminal records of its registered users. Thus, there is no need to learn for which particular user a request is made, or whether several requests belong together. In our system, the converter is able to provide such access control but does not learn any additional information from the queries.

We formally define the functional and security properties such an (un)linkable pseudonym system should ideally provide. Our security definition is formulated in the universal composability framework, and thus comes with strong guarantees when composed with other UC secure protocols. We then describe our system using generic building blocks.

The idea of our solution to build pseudonyms by adding several layers of randomness to the user identifier, such that they allow for consistent (and blind) conversions yet hide the contained identifier towards the servers. Roughly, to generate a pseudonym $nym_{i,A}$ for a user uid_i on server $\mathcal{S}_A$, the converter first applies a verifiable PRF on uid_i and then raises the derived value to a secret exponent that it assigns for each server. The trick thereby is that those secret keys are known only to the converter, but are never revealed to the servers.

Now, consider the blind conversion procedure. It can of course be realized with a generic multiparty protocol, where the first server $\mathcal{S}_A$ inputs the pseudonym to be converted and the converter inputs all its secret keys, and the output of the second server $\mathcal{S}_B$ would be the converted pseudonym, provided that the input by $\mathcal{S}_A$ was a indeed a valid pseudonym. However, such a computation would be rather inefficient. We therefore aim to construct a specific protocol that achieves this efficiently.

We propose a blind conversion protocol that performs the conversion on *encrypted* pseudonyms, using a homomorphic encryption scheme. To transform a pseudonym from one server to another, the converter then exponentiates the encrypted pseudonym with the quotient of the secret keys of the two servers. The challenge is to make that entire process verifiable, ensuring that the conversion is done in a consistent way but without harming the privacy properties.

To ensure controllability in the sense that a server can only request conversions for pseudonyms it legitimately obtained via the converter, we also make use of a novel building block which we call *dual-mode signatures*. Those allow to obtain signatures on encrypted messages, which can then be "decrypted" to a signature on the underlying plaintext message. We also provide a concrete construction for those signatures based on the recent signature scheme by Abe et al. [16], which might be of independent interest. Our dual-mode signatures can be seen as a specialised variant of commuting signatures [182], and therefore allow for more efficient schemes.

Finally, we prove that our protocol realizes our ideal functionality based on the security of the building blocks. We also provide concrete instantiations for all generic building blocks used in our construction which already come with optimizations and enhance the efficiency of our solution. When instantiated with the suggested primitives, our protocol is secure based on discrete-logarithm related assumptions.

7 Research on Privacy-Friendly Revocation Mechanisms

In Section 7.1, we propose a privacy preserving revocation mechanism for privacy-enhancing attribute-based credentials that allows you to efficiently handle multiple revocation lists. In Section 7.2, we study a primitive that is widely used for revocation purposes, i.e., cryptographic accumulators. In Section 7.3 we show how using epochs can help to make revocation practical while still retaining reasonable strong privacy guarantees. Our contribution is a new revocation scheme that has very low computational cost for users and verifiers alike, that is efficient even in the smart card setting, and therefore can be used in practice. In Section 7.4 we explain the concept of revocable privacy.

7.1 UC Commitments, Revocation, and Attribute Tokens for Privacy Preserving Protocol Design

Privacy-enhancing attribute-based credentials (PABC) [109, 113], also known as anonymous credentials [142, 116, 114, 117, 46, 47, 99, 40] or minimal-disclosure tokens [77], are cryptographic mechanisms to perform data-minimizing authentication. Contrary to other attribute certification schemes (e.g., SAML, OpenID, X.509 certificates), PABCs offer important privacy advantages, such as unlinkability, selective disclosure and proving in zero-knowledge predicates over the attributes [109]. Traditional certificate revocation mechanisms do not apply to PABCs because they do not preserve unlinkability. Therefore, different privacy-preserving revocation mechanisms have been proposed in the literature such as signature lists [272], accumulators [115, 278, 104], and validity refreshing [105].

All the existing revocation mechanisms assign a binary value to each credential, specifying whether the credential is valid or revoked. However, that approach does not effectively address all the many different reasons for which credentials should be revoked. In some cases, credentials need to be revoked globally, e.g., when the related secret keys are exposed, the attribute values have changed, or the user loses her right to use a credential. Often, credentials may be revoked only for specific contexts, i.e., when a user is not allowed to use her credential with a particular verifier, but can still use it elsewhere. For example, a stadium may require spectators to prove possession of a valid identity card and prevent blacklisted spectators from accessing the stadium. However, they can still use their identity card for other purposes. Another scenario is a dynamic attribute-based access control system, where the users are issued several credentials of different types, each certifying a certain attribute or role, but bound to the same user secret. If a certain attribute or a role needs to be revoked, the corresponding credential gets revoked.

In such scenarios, the revocation authority needs to maintain multiple revocation lists. Because of their binary value limitation, the existing revocation systems require applying a revocation mechanism for each list separately. This imposes an extra storage and computational overhead not only to the users, but also to the revocation authority. Furthermore, in signature lists and accumulators, the revocation lists are disclosed to the other users and verifiers. Therefore, creating a privacy-preserving revocation scheme that supports hidden revocation statuses and multiple revocation lists efficiently still remains an important open problem.

Privacy and efficiency are not the only concerns when building complex cryptographic protocols, but the protocols need also be proven secure and retain their security when composed with other protocols. Unfortunately, many schemes are still proven secure in standalone models that guarantee security only if no other protocols are run concurrently, which is not the case for a real environment. Moreover, the security proofs are monolithic and therefore very complex, prone to mistakes and hard to verify. Composability frameworks such as the UC framework [128] guarantee that security is preserved even if the protocols are used as components of a complex system, enabling a modular design and security analysis of cryptographic protocols. Thus, when constructing a complex credential system, each building block by itself should provide certain security and privacy guarantees and be easily composable with the rest of the system.

The advanced features and security requirements of PABC result in complex systems and make it a prime target for modular design and security proofs. The difficulty in achieving this is that composing different ideal functionalities to build the overall systems requires that cryptographic values be exchanged between these functionalities. Such cryptographic values include commitments or signatures. For example, in order to verify a proof that the credential was not revoked and evaluate a predicate over the attributes from this credential using different building blocks, one needs to be sure that different parts of the proof are done with respect to the same credential, bound to the same user secret. Another example is when using commuting signatures [182] that, given a verification key, a message and a signature on it valid under the key, allow one to encrypt (commit to) any subset of the key, message, or signature, and prove that the plaintexts constitute a triple of a key, a message, and a valid signature. Unfortunately, very few ideal functionalities found in the literature output cryptographic values, and, if they do, it is not possible to use these values in another functionality in an essential way (apart from treating them as an opaque blob). For instance, while commitment schemes are a perfect tool for bridging different building blocks in a privacy-preserving manner and have been used as such in the design of many cryptographic protocols, the existing ideal functionalities for commitments do not output cryptographic values. Thus is it not possible to use them as a building block for protocols whose security can be proven in a modular way – defeating the whole purpose of composability frameworks. Therefore, defining such functionalities and instantiating them is another crucial task that needs to be completed in order to build composable privacy-preserving systems from different cryptographic building blocks, such as signatures, verifiable encryption, accumulators, etc.

We address all the issues described above and present a number of new UC functionalities, provide efficient realizations and prove them secure, and finally give an example of how to construct and prove secure a high-level protocol in a hybrid model using our new functionalities. The building blocks comprise flexible revocation mechanisms, trapdoor vector commitments, and a bare-bones credential scheme to certify attributes. As an example high-level scheme in a hybrid model, we construct a credential scheme with revocation. We discuss each of the contributions in more detail below.

Multiple revocation lists. We propose a mechanism that allows one to accumulate several revocation lists into a single commitment value, each of the lists containing the valid revocation handles (e.g., serial numbers of the related credentials). Each user needs only one witness for

all the multiple revocation lists. Using this witness a user can prove in a privacy-preserving manner whether or not her revocation handle is on a particular revocation list. We provide the ideal functionality for such mechanisms and propose two different constructions. The first one hides the revocation status of a user from other users and from the verifiers while in the second one, as for accumulators, revocation lists are public. Additionally, our schemes are flexible in the sense that revocation lists can be added (up to a maximum) and removed without any cost, i.e., the cost is the same as for a revocation status update that does not change the number of lists, while accumulators would require setting up a new accumulator and issuing witnesses to users, or deleting those.

We note that aside from extending the standard revocation scenario with a central revocation authority and multiple revocation lists, our revocation schemes can be used to build an efficient dynamic attribute-based access control system in a very elegant way. Instead of issuing to each user a list of credentials, each certifying a certain attribute or role, using our revocation schemes a user can be issued just one base credential, which can be made valid or revoked for all the different contexts. The resulting solution saves the users, verifiers and the revocation authority a lot of storage and computational effort. That is, instead of having multiple credentials and corresponding revocation witnesses, just a single credential and a single witness suffice to achieve the same goal.

Trapdoor vector commitments. We build both of our revocation schemes from vector commitments [132]. Vector commitments allow one to commit to a vector of messages and to open the commitment to one of the messages in such a way that the size of the opening is independent of the length of the vector. We employ trapdoor and non-hiding vector commitments for our first and second construction, respectively. To prove our non-hiding construction secure we rely on the binding property for non-hiding vector commitments [132]. For the hiding construction, we define a trapdoor property that allows the simulator to open a vector component to any value, while still ensuring indistinguishability between real and fake openings.

Ideal functionalities. To address the composability issue we define the ideal revocation functionalities for both hiding and non-hiding options in the universal composability framework [128]. To make our revocation building blocks employable by other protocols, such as attribute-based credentials systems, we implement the following features. First, we let the parameters of the revocation scheme depend on parameters generated externally. These parameters also include the commitment parameters and a commitment verification algorithm. This allows the revocation functionality to compute and verify revocation proofs about commitments to revocation handles, such that the commitments are generated externally, e.g., by the commitment functionality. In this case, the commitment functionality bridges revocation proofs with other proofs, for example, that the credential indeed contains this revocation handle. This requires the commitment functionality to output a cryptographic value after receiving a commit message.

To achieve that, we provide a new ideal functionality of commitment schemes that can be easily employed as a building block to create complex systems. We show that any trapdoor commitment scheme realizes our functionality. We believe that our work on the trapdoor com-

mitment and vector trapdoor commitment schemes is of independent interest and can be used outside of the credentials context. We show that, surprisingly, the trapdoor property suffices to prove the system composably secure. Intuitively, one might consider simulation-sound trapdoor commitments [252]. However, this property turns out to be too strong and not needed. A simulation-sound property for vector commitments is also unnecessary.

As another building block, we give the first UC functionality and realization of an (anonymous) attribute token scheme [121].

Composable revocable tokens. Finally, we demonstrate the feasibility of our approach and build a variant of the credential system from the proposed building blocks. Namely, we construct universally composable anonymous attribute tokens that support revocation. However, using our building blocks one can create complex composable systems, such as composable commuting signatures, anonymous credential systems, etc.

7.2 Revisiting Cryptographic Accumulators, Additional Properties and Relations to other Primitives

A (static) cryptographic accumulator scheme allows to accumulate a finite set $\mathcal{X} = \{x_1, \ldots, x_n\}$ into a succinct value $\mathsf{acc}_{\mathcal{X}}$, the so called accumulator. For every element $x_i \in \mathcal{X}$, one can efficiently compute a so called witness $\mathsf{wit}x_i$ to certify the membership of x_i in $\mathsf{acc}_{\mathcal{X}}$. However, it should be computationally infeasible to find a witness for any non-accumulated value $y \notin \mathcal{X}$ (*collision freeness*). Dynamic accumulators are an extension that allows to dynamically add/delete values to/from a given accumulator and to update existing witnesses accordingly (without the need to fully recompute these values on each change of the accumulated set). Besides providing membership witnesses, universal accumulators also support non-membership witnesses for values $y \notin \mathcal{X}$. Here, collision freeness also covers that it is computationally infeasible to create non-membership witnesses for values $x_i \in \mathcal{X}$. Over time, further security properties, that is, *undeniability* and *indistinguishability* have been proposed. Undeniability is specific to universal accumulators and says that it should be computationally infeasible to compute two contradicting witnesses for $z \in \mathcal{X}$ and $z \notin \mathcal{X}$. Indistinguishability says that neither the accumulator nor the witnesses leak information about the accumulated set $\mathcal{X}$ and, thus, requires randomized accumulator schemes.

Applications: Accumulators were originally proposed for timestamping purposes [55], i.e., to record the existence of a value at a particular point in time. Over time, other applications such as membership testing, distributed signatures, accountable certificate management [85] and authenticated dictionaries [193] have been proposed. Accumulators are also used as building block in redactable [296, 297], sanitizable [126], P-homomorphic signatures [23], anonymous credentials [325], group signatures [333], privacy-preserving data outsourcing [319] as well as for authenticated data structures [191]. Moreover, accumulator schemes that allow to prove the knowledge of a (non-membership) witness for an unrevealed value in zero-knowledge (introduced for off-line e-cash in [314]) are now widely used for revocation of group signatures and

anonymous credentials [115]. Quite recently, accumulators were also used in Zerocoin [261], an anonymity extension to the Bitcoin cryptocurrency. Since their introduction, numerous accumulator schemes with somewhat different features have been proposed. Basically, the major lines of work are schemes in hidden order groups (RSA), known order groups (DL) and hash-based constructions (which may use, but typically do not require number theoretic assumptions).

Hidden order groups: The original RSA-based scheme of Benaloh and de Mare [55] has been refined by Baric and Pfitzmann [43], who strengthen the original security notion to *collision freeness*. In [313], Sander proposed to use RSA moduli with unknown factorization to construct trapdoor-free accumulators. Camenisch and Lysyanskaya [115] extended the scheme in [43] with capabilities to dynamically add/delete values to/from the accumulator, which constituted the first *dynamic accumulator* scheme. Their scheme also supports *public updates* of existing witnesses, that is, updates without the knowledge of any trapdoor. Later, Li et al. [240] added support for non-membership witnesses to [115] and, therefore, obtained *universal dynamic accumulators*. They also proposed an optimization for more efficient updates of non-membership witnesses, for which, however, weaknesses have been identified later [294, 255]. Lipmaa [245] generalized RSA accumulators to modules over Euclidean rings. In all aforementioned schemes, the accumulation domain is restricted to primes in order to guarantee collision freeness. In [333], Tsudik and Xu proposed a variation of [115], which allows to accumulate semiprimes. This yields a collision-free accumulator under the assumption that the used semiprimes are hard to factor and their factorization is not publicly known. Moreover, in [337] an accumulator scheme that allows to accumulate arbitrary integers and supports batch updates of witnesses has been proposed. Yet, this scheme was broken in [88].

Known order groups: In [278], Nguyen proposed a dynamic accumulator scheme which works in pairing-friendly groups of prime order p. It is secure under the t-SDH assumption and allows to accumulate up to t values from the domain $\mathbb{Z}_p$. Later, Damgård and Triandopoulos [157] as well as Au et al. [35] extended Nguyen's scheme with universal features. Quite recently, Acar and Nguyen [20] eliminated the upper bound t on the number of accumulated elements of the t-SDH accumulator. To this end, they use a set of accumulators, each containing a subset of the whole set to be accumulated. An alternative accumulator scheme for pairing friendly groups of prime order has been introduced by Camenisch et al. [104]. It supports public updates of witnesses and the accumulator and its security relies on the t-DHE assumption.

Hash-based constructions: Buldas et al. [85, 86] presented the very first universal dynamic accumulator that satisfies *undeniability* (termed as undeniable attester and formalized in context of accumulators in [245]). Their construction is based on collision-resistant hashing and the use of hash-trees. Another hash-tree based construction of a universal accumulator that satisfies a notion similar to undeniability has been proposed in [89] (the scheme is called a strong universal accumulator). Quite recently, another accumulator based on hash-trees, which uses commitments based on bivariate polynomials modulo RSA composites as a collision-resistant hash function, has been introduced in [66]. For the sake of completeness, we also mention the

construction of static accumulators in the random oracle model based on Bloom filters, proposed by Nyberg [284, 283].

We address the following issues.

Unified model: While some papers [55, 43, 115, 35, 278] do not explicitly formalize accumulator schemes, formal definitions are given in [104, 337, 174, 240, 132, 245, 89, 20]. However, these models are typically tailored to the functionalities of the respective scheme. While they widely match for the basic notion of (static) accumulators (with the exception of considering randomized accumulators), they differ when it comes to dynamic and universal accumulators. To overcome this issue, we propose a unified formal model for accumulators, which is especially valuable when treating accumulators in a black-box fashion. We, thereby, also include the notion of undeniability [85, 86, 245] and a strengthened version of the recent indistinguishability notion [159]. Besides, we also confirm the intuition and show that undeniability is a strictly stronger notion than collision freeness.

Classification: We provide an exhaustive classification of existing accumulator schemes and show that most existing accumulator schemes are distinguishable in our model. To resolve this issue, we propose a simple, light-weight generic transformation that allows to add indistinguishability to existing dynamic accumulators and prove the security of the so-obtained schemes. As this transformation, however, comes at the cost of reduced collision freeness, we additionally propose the first indistinguishable scheme that does not suffer from this shortcoming.

Relations to other primitives: Since accumulators are somehow related to commitments to sets [228, 173], commitments to vectors [132] and to zero-knowledge sets [259], it is interesting to study their relationship. Interestingly, we can formally show that indistinguishable accumulators imply non-interactive commitment schemes. Furthermore, we formally show that zero-knowledge sets imply indistinguishable, undeniable universal accumulators, yielding the first construction of such accumulators.

7.3 Fast Revocation of Attribute-Based Credentials for Both Users and Verifiers

Governments increasingly issue electronic identity (eID) cards to their citizens [236, 275, 285]. These eID cards can be used both offline and online for secure authentication with the government and sometimes with other parties, like shops. Attribute-based credentials (ABCs) [108] are an emerging technology for implementing eID cards because of their flexibility and strong privacy guarantees, and because they can be fully implemented on smart cards [336]. Every credential contains attributes that the user can either reveal or keep hidden. Such attributes describe properties of a person, like her name and age. ABCs enable a range of scenarios from

fully-identifying to fully-anonymous.[5] When using a credential fully anonymously (i.e., without revealing any identifying attributes), proper ABC technologies guarantee that the credential is unlinkable: it is not possible to connect multiple usage instances of the same credential.

When ABCs are applied, the carriers on which the credentials are stored (for example smart cards) can be lost or stolen. In such cases it is important that users can revoke these credentials, to ensure that they can no longer be (ab)used. This is also necessary when the owner of the credential herself abuses it. Revocation may in fact happen often. As an example, the nationwide Belgian eID system's revocation list contains more than 375 000 credentials [104] for just over 10 million citizens. A practical revocation scheme must therefore efficiently deal with such large revocation lists.

Unfortunately, the unlinkability of ABCs precludes the use of standard, identity-based, revocation. There exist many privacy-friendly revocation schemes, with different trade-offs in terms of efficiency (both for users and verifiers), connectivity requirements, and anonymity. It turns out to be hard to satisfy all of these simultaneously. In particular, all revocation schemes proposed so far suffer from at least one of the following two problems: (1) they rely on computationally powerful users, making the scheme unsuitable for smart cards, the obvious carrier for a national eID card; or (2) they place a high load on verifiers, resulting in long transaction times.

Our contribution. Our contribution is a new revocation scheme that has very low computational cost for users and verifiers alike, that is efficient even in the smart card setting, and therefore can be used in practice. In our scheme, verifiers need only constant time on average to check revocation status, making it as fast as traditional non-anonymous revocation schemes. Furthermore, the users' computational overhead is small (and updates to reflect new revocations are *not* necessary). Our scheme is unlinkable, except if the user uses her credential more than once per epoch at the same verifier.

Sketch of the scheme Our scheme resembles verifier-local revocation (VLR) schemes [33, 71, 81], and is based on *epochs* that divide time in short (system-configurable) intervals. The essence of VLR schemes is that a verifier receives a list of blacklisted credentials. These blacklists contain, for each revoked credential, a value g^r for some common base point g and a revocation value r that is also embedded in the credential. When a prover needs to prove possession of this credential, it is challenged with the base point g. It needs to compute g^r using the revocation value r embedded in the credential, and prove that the value r he used corresponds to the value in the credential. Clearly, the prover is traceable because g^r is fixed. In our approach this is fixed in two ways. First, g is different for each verifier. This means provers are no longer traceable across verifiers. Second, g depends on the current epoch. This ensures that provers are no longer traceable for some length of time across epochs.

[5] This is why we prefer the term 'attribute-based credentials' over the more traditional term 'anonymous credentials'.

7.4 Revocable Privacy: Principles, Use Cases, and Technologies

Privacy and (homeland) security seem to be at odds with one another: it is a commonly held belief that we cannot strengthen one without weakening the other. And it seems security is winning. The governmental hunger for data—and its ability to actually gather these—seems bigger than ever. And who would argue against collection of these data? Surely we all want to stop terrorists, paedophiles and tax evaders. Yet, security versus privacy does not have to be a zero-sum game [315, 323]. Hoepman also argued that this contradiction between security and privacy is a false one, and that we can design systems that have privacy without neglecting security [215].

Hoepman introduced a design principle to create systems that have both security and privacy: *revocable privacy*. The core idea of revocable privacy arises from the realisation that it is not the data itself that is (or should be) important, but rather the violations of certain rules that manifest themselves in the data. Data related to people who do not violate any rules are irrelevant, and, in fact, these people should remain anonymous, as if no data on their behavior was ever collected. Revocable privacy is a design principle that ensures this property. Informally speaking, a system offers revocable privacy if users of the system are guaranteed to be anonymous except when they violate a predefined rule.

To ensure privacy, the system's anonymity guarantees cannot rely on policy and regulations alone. It is all too easy to ignore policy, to sidestep it, or to change it retroactively. As a result, data that was collected for one specific purpose can easily be reused for another—violating people's privacy. A key aspect of a system implementing revocable privacy is to prevent this type of function creep through technical means: it should not be possible to change the rules retroactively.

It is known that building such systems is possible. One example is the anonymous electronic cash system proposed by Chaum [140], which actually implements revocable privacy (although he did not use this term). Users have electronic coins, which they can spend as if they were physical coins, in effect making an untraceable digital payment system in which the users' privacy is guaranteed. However, to maintain security, this anonymity cannot be unconditional. If it were, it would allow misbehaving users to double-spend the digital coins without consequence. Instead, the revocable privacy aspect of the design guarantees that users are anonymous, as long as they spend the digital coins only once. When they do spend a coin twice, their identity can be recovered from the two transaction records of the two spendings. Any single transaction record, however, gives no information about the identity of the user.

In general, to ensure anonymity for rule-abiding users, data must be collected in a special manner. In Chaum's electronic cash system, the cut-and-choose paradigm is used to ensure that a single transaction record gives no information, whereas two reveal the identity of the culprit. Distributed encryption [216, 249] offers another method for creating threshold based rules. Data is collected for every event, but the user's identity is revealed only if she causes an event to happen at sufficiently many different locations.

While Chaum's electronic cash could be seen as such a scheme with a threshold of two, it differs significantly from distributed encryption. In the first, the user actively partakes in the

transaction, whereas in the second, the user deliberately does not take part. As a result, these systems have different privacy guarantees and trust assumptions. These aspects of revocable privacy had not yet been explored.

In all previous work on revocable privacy [215, 216, 248, 249], the focus was on identifying users who violate the rules. However, in some situations, such an approach might be too strong. For example, anonymity is the core property of Tor [167], so it should never be possible to deanonimize users. Yet, Tor can also be abused. In order to stop abuse, some approaches, like blacklistable anonymous credentials (BLAC), aim to block misbehaving users, rather than to identify them [331].[6]

Our first contribution is to reexamine revocable privacy in a more general setting, where we consider the implications of different security models, and explore ramifications of users' actions that are less invasive than simply identifying users, for example, blocking users and linking their actions.

Next, we explore and classify some use cases for revocable privacy. We generalize the underlying rules of the use cases into abstract rules. These use cases illustrate that even if a user has violated a rule, she did not necessarily do something wrong. In fact, we have explored some systems where a violation only means that closer examination is necessary.

The abstract rules for the use cases make it possible to link them to specific techniques. Our final contribution is to give a non-technical overview of existing techniques that can be used to implement revocable privacy.

[6] Nymble [332] is a related system that can be used to block misbehaving users. However, it relies on a trusted party that can make users linkable if they misbehave, so we do not consider it further in this paper.

8 Methods for Usable Privacy

In Section 8.1, we present a two-factor user-authentication scheme for usable server-based eID and e-signature solutions. Current server-based eID and e-signature solutions typically rely on one-time passwords delivered to the user via short message service (SMS). This raises several issues in practice, as the use of SMS technology can be cost-effective insecure. To address these issues, we propose an alternative two-factor user-authentication scheme following a challenge-response approach. The feasibility and applicability of the proposed user-authentication scheme is evaluated by means of two concrete implementations. This way, we show that the proposed authentication scheme and its implementations improve both the cost effectiveness and the security of server-based eID and e-signature solutions.

In Section 8.2, we study how users choose passwords under consideration of different human dimensions, and, more specifically, when they are cognitively depleted.

8.1 Encryption-based Second Authentication Factor Solutions for Qualified Server-side Signature Creation

The concepts of electronic identity (eID) and electronic signature (e-signature) are crucial for transactional e-government services. They enable users to securely and reliably authenticate at services and to create electronic signatures. Their relevance is especially given in the European Union (EU), where so-called qualified electronic signatures are legally equivalent to handwritten signatures [328]. This enables users to remotely provide written consent in transactional services.

During the past years, different approaches for the realization of eID and e-signature concepts have been studied, implemented, and deployed. First approaches to provide users eID and e-signature functionality have been based on smart-card technology [237, 172]. However, these approaches have turned out to suffer from several usability-related limitations and hence from limited user acceptance [347]. As an alternative, mobile eID and e-signature solutions have emerged early. These solution remove the need for smart-card usage by making use of the user's mobile phone instead. Two approaches can be distinguished. The first approach employs the mobile phone's SIM card to store eID data and to implement cryptographic functions required for the creation of electronic signatures. The second approach instead relies on a central hardware security module (HSM) to store eID data and to carry out required cryptographic functions.

During the past years, the second approach, i.e., server-based solutions, has gained relevance and popularity, mainly because it defines fewer requirements for the mobile end-user device and mobile network operators (MNO), which in turn improves applicability, feasibility and usability. The main challenge in designing and developing server-based eID and e-signature solutions is the provision of appropriately secure user-authentication schemes. These schemes are required to restrict access to centrally stored eID data and cryptographic signing keys to the legitimate user. In order to assure a sufficient level of security, two-factor authentication (2FA) is typically the approach of choice.

Current mobile eID and e-signature solutions following the sever-based approach implement 2FA

schemes by means of one-time passwords (OTPs) delivered by SMS messages [290, 303]. After the user has entered a secret password covering the authentication factor knowledge, he or she receives a OTP via SMS. By proving reception of the OTP, the user proves possession of the mobile phone. This way, the authentication factor possession is covered and the 2FA process is completed.

Unfortunately, reliance on SMS technology raises several issues [271]. First, SMS must not be regarded as secure. This especially applies to smartphones, on which incoming SMS messages can be intercepted by third party applications. Second, the sending of SMS messages containing OTPs can cause significant costs for the service operator, as mobile network operators (MNOs) typically charge the delivery of SMS messages. To overcome these issues, we propose an alternative 2FA scheme for server-based mobile signature solutions. Our proposed scheme renders the use of SMS technology unnecessary. This way, it provides higher cost efficiency and better security compared to existing approaches.

8.2 Human Dimensions of Identity Federation and Password Choice

In earlier research [151], we proposed and subsequently investigated the impact of human factors, such as cognitive effort, on privacy decision making. This research is based on the observation in psychology research that cognitive effort can be measured with task-evoked pupillary response [227, 225] and that cognition can be viewed in a dual-process model that distinguishes cognitively effortless and effortful tasks [226]. Naturally, the question arose whether security and privacy decisions are inherently effortful or inherently effortless. Furthermore, it is interesting whether users whose capacity for effortful decisions is depleted make security and privacy decisions differently.

Research Methodology for Human Dimensions in Security and Privacy Such experiments require an accurate and calibrated measurement apparatus and a rigorous evidence-based research methodology. We also found that such experiments come with a number of confounding variables that need to be tightly controlled [150]. This prompted the need for high-end psycho-physiological measurements as a tool for further investigation. Newcastle University has designed and built a psycho-physiological measurement and eye tracking lab to enable this research. It includes a high-end eye tracker (SMI RED500), a system to measure electro cardiograms (ECG) and electrodermal activity (EDA), a system for affect analysis from face-geometry as well as an analysis system to investigate these multimodal data streams and to add further coding of user behavior. The bottomline is that one can measure cognitive and affective factors while the user is working with a computer making security and privacy decisions.

Current research includes validating the psycho-physiological measurements against established psychological tools and defining reusable components for future experiments. For instance, a current experiment validates the affect evaluation from face-geometry (Noldus FaceReader) against a standard questionnaire for positive and negative, the Positive and Negative Affect Schedule (PANAS-X) [339]. Furthermore, this research requires standard methods for manipulating the user's cognitive and affective state, such as inducing cognitive effort or depletion.

Human Dimensions of Password Choice A first larger experiment with this methodology focused on a simple question that is closely related to identity federation: How do users choose passwords under consideration of different human dimensions? More specifically how do users choose passwords when they are cognitively depleted? This question is important for FutureID, as passwords are still a predominant authentication mechanism for Identity Providers (IdPs). We summarize this experiment to illustrate a research methodology and study developed in FutureID which can be used for further investigations in Usable Privacy. The current study investigated the password strength and memorability, measured in the outcome of a password meter and login attempts a week after registration.

Method. In this still unpublished study, 100 participants were chosen and randomly distributed across two groups of 50 participants each. The groups were balanced in terms of gender and time-of-day to run the experiment to for control circadian rhythm. We queried demographics, gender, a security awareness score and a Big Five personality inventory (BFI) [223] as pre-study questionnaire. The participants were mostly non-computer science students from Newcastle and Northumbria University, mostly international background (most common countries Oman, China, Iraq), with ages between 17 and 44 years.

The experiment group was artificially cognitively depleted with tasks that required impulse control, the control group was not depleted, completing non-depleting tasks with similar length and flavor.

The password strength was measured with a password meter[7] and adjusted for the use of unmodified dictionary words and personal identifiable information in the password (e.g., username or student id) as per NIST recommendations on good password choice.

Manipulation Check. A manipulation check was employed with a brief mood inventory used in earlier psychology research, yielding that the manipulation was successful with statistical significance (Mann-Whitney, significance $p = .000 < .05$). From the manipulation check, a depletion level was derived with levels of non-depleted, cognitively effortful, and depleted. A total of 28 participants reported a level of depletion expressed as slight or strong tiredness. Participants reporting neither agreement or disagreement to tiredness were counted as non-depleted.

Results. The data was analyzed with a multi-predictor stepwise linear regression. The linear regression accounted for 20.6% of the variability (adjusted $R^2 = .206$), hence offers only a limited accuracy. The studentized residual was close to a normal distribution. Nearly 80% of the variability was unaccounted for and there was a large variance observed within groups.

The outcomes of the gender, Big Five (5), brief mood inventory (8) and the depletion level were predictors on the password strength score as target variable. The depletion level was indeed the most important predictor in the regression (significance $p = .001 < .05$, predictor importance= 0.371). The effortful level, that is only slightly depleted, had a coefficient of 50.65 (significance $p = .000 < .05$). The non-depleted level had a coefficient of 31.62 (significance $p = .006 < .05$). We summarize these descriptive statistics of password strength score by depletion level in Table 2 and depict the means with 95%-confidence intervals in Figure 1.

[7] `www.passwordmeter.net`

Table 2: Descriptive statistics of password strength score by depletion level.

Depletion Level	N	Mean	Std. Dev.	Std. Error	Min	Max
Non-depleted	72	38.31	31.86	3.76	-45	121
Effortful	17	57.12	45.07	10.93	-24	138
Depleted	11	11.55	43.63	13.16	-64	70
Total	100	38.56	37.27	3.73	-64	138

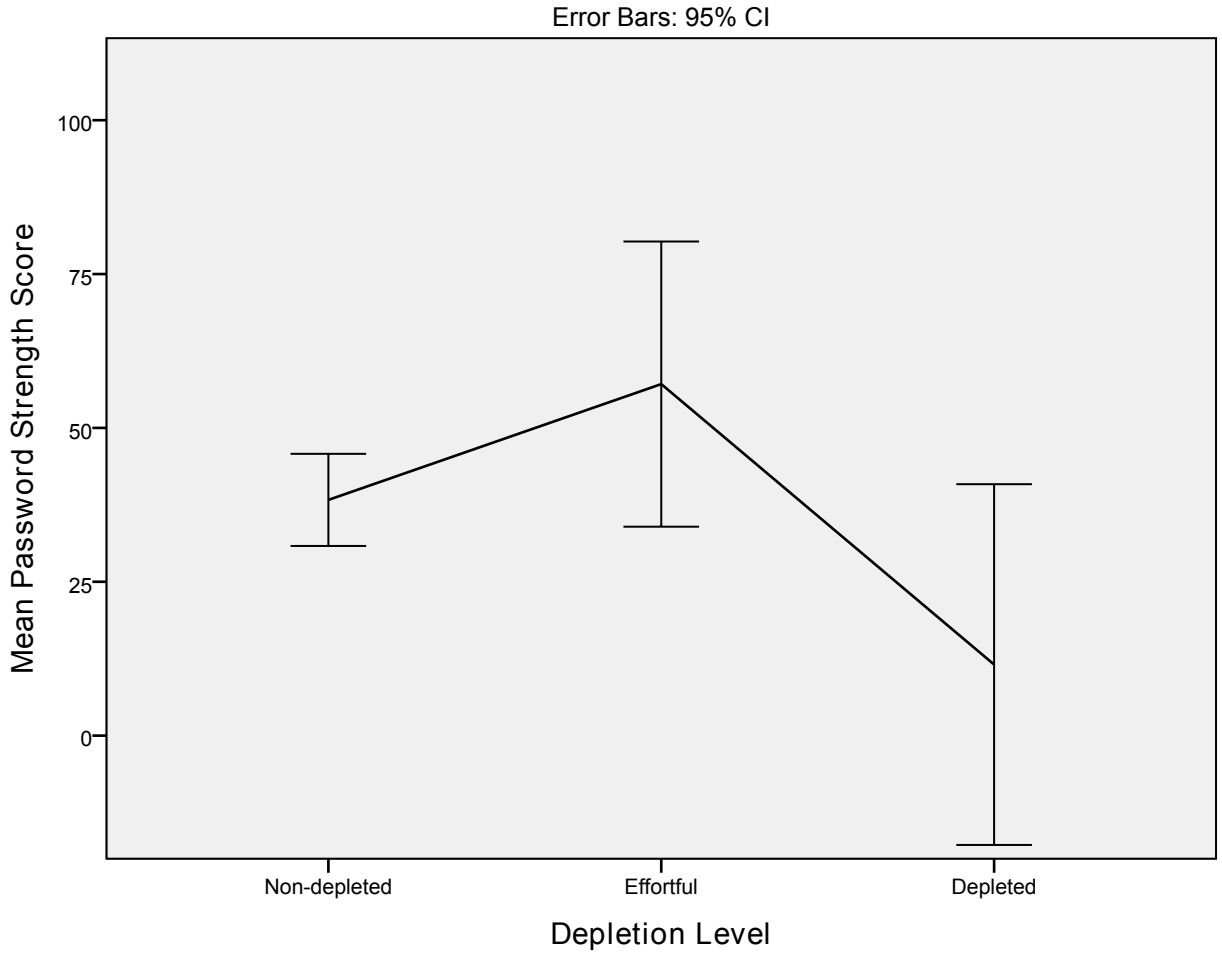

Figure 1: Means of password strength score by depletion level.

Of the brief mood inventory thoughtfulness and calmness had significant effects. Strong disagreement to thoughtfulness implied stronger passwords (significance $p = .018 < .05$, predictor importance $= 0.251$). Strong disagreement to calmness implied stronger passwords (significance $p = .012 < .05$, predictor importance $= 0.172$).

Of the Big Five personality traits, the BFI Agreeableness score was the most important predictor on the password strength (significance $p = .025 < .05$, predictor importance $= 0.137$), where higher agreeableness significantly implied stronger passwords. The BFI Extraversion was a notable yet non-significant negative predictor on password strength (significance $p = .108 > .05$, predictor importance $= 0.069$).

Discussion. This outcome on depletion level as major predictor suggests that slight cognitive effort prior to a password choice implies a stronger password chosen. It also suggests that a strong depletion implies a weaker password chosen. This is in accord with Kahneman's observation that initial effortful activity introduces a bias towards exerting further cognitive effort [226]. This outcome can also be explained with Selye's arousal curve [146], an inverse U-shaped relation between the activity of the stress system and the quality of a human's performance, yielding an optimum performance under moderate stress. This result vouches for further investigation, in particular to what extent this observation can be operationalized to improve the quality of password choice.

The results on the brief mood inventory are surprising in themselves. Why do participants who report themselves as not thoughtful or not calm choose better passwords? This result can substantiate the explanation of Selye's arousal curve as a possible explanation. In any case, these results ask for the investigation of the influence of current stress and affect on password choice.

This result on the BFI asks for further investigation as the experiment cannot explain whether the participants sought to please the experimenter, constituting a confounding variable, or whether this effect persists in real-world scenarios.

The number of participants that reported strong depletion was low ($N = 11$), limiting the importance of this coefficient. Future Experiments will need to achieve stronger depletion throughout and manipulate a slight cognitive effort stimulus deliberately.

Conclusion We believe that the research methodology developed in FutureID for the evaluation of human dimensions of security and privacy will be widely reusable. They can shed light on human dimensions of identity federation beyond the capabilities of existing self-report questionnaires on usability. Hence, this work constitutes a stepping stone for further research in usable privacy and the security of identity federation systems. We advocate future research using psycho-physiological measurements and differentiated consideration of personality traits and user experience as a way forward to gain more insights in usable security and privacy.

9 Conclusion

This deliverable has described the research conducted by work package 24 during the third 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).

In Task 24.1, we have achieved a significant step forward in the area of compositional reasoning with two kinds of *relative soundness results*. The first kind are typing results showing that any security protocol that fulfils a number of sufficient conditions has an attack if it has a well-typed attack. The second kind considers the parallel composition of protocols, showing that when running two protocols in parallel allows for an attack, then at least one of the protocols has an attack in isolation.

In Task 24.2, we have focused on applying the concept of α-β-privacy that we have developed previously to the FutureID architecture. Furthermore, we have defined and unified the concepts and features of privacy-preserving attribute-based credentials (Privacy-ABCs), provided a language framework in XML schema, and given a formal semantics to describe the effects of the transactions in a privacy-friendly authentication system using Privacy-ABCs. Additionally, we present a Prolog implementation for credential-policy matching.

In Task 24.3, for audits, we have reported on experiences during implementing blank digital signatures as well as optimizations that helped to improve their performance. Additionally, we have proposed a novel graph signature scheme, which makes it possible that an issuer certifies a committed graph, such that a prover can subsequently prove properties of the graph in zero-knowledge proofs of knowledge.

For data-handling mechanisms, we have focused our research on authentication mechanisms based on passwords and on signatures and on privacy-preserving protocols that minimize the data that users have to disclose to service providers. We have also worked on existing eID solutions. In addition, we have conducted research on computations on signed data and on data anonymization and data sharing between databases.

Password-Based Authentication. We have described a threshold password authenticated secret sharing protocol. We have also presented two simple and extremely efficient proactively secure *distributed password verification* protocols.

Signature-Based Authentication. We have given formal security definitions for a full-fledged privacy attribute-based credentials system. We have provided a generic construction from lower-level building blocks that satisfies our definitions and we have presented secure instantiations of the building blocks. Additionally, we have proposed a new kind of signature schemes, *unlinkable redactable block-signature* (URS) schemes, with which one can redact a signature and reveal only its relevant parts each time it is used. We have constructed an efficient URS scheme and we have employed it to design the first universally composable anonymous credential system. It is also arguably one of the first such schemes to support

efficient attribute disclosure with cost independent of the number of attributes in the issued credential without having to rely on random oracles. Moreover, we have presented an efficient construction of round-optimal blind signature schemes in the standard model.

eID Solutions. We have presented design strategies for a privacy-friendly Austrian eID system in the public cloud.

Privacy-Preserving Protocols. We have contributed a practical, secure and privacy preserving mechanism enabling a Service Provider to verify whether the mobile phone of a given User currently resides within a certain geographical reference area at a given time. Our mechanism consists of having the location of the mobile phone determined by the Mobile Network Operator and certified using anonymous credentials. In addition, we have proposed a non-restricted and a restricted oblivious transfer with access control scheme. Oblivious transfer with access control (OTAC) allows the sender to control access to the messages. The sender receives as input a list of messages and access control policies. Each receiver possesses a set of attributes, which is certified by a credential issuer, and is able to obtain a message only if the receiver's attributes satisfy the access control policy for that message. Receiver privacy requires that the sender does not get any information on the message obtained or on the receiver's attributes. In a non-restricted scheme, a receiver can obtain in one transfer phase all the messages whose access control policy is fulfilled by the receiver's attributes. In a restricted scheme, the receiver can only obtain one message per transfer phase. Furthermore, we have revisited existing work on privacy-preserving billing. In privacy-preserving billing a meter measures a user's consumption of some utility or service and service providers apply fine-grained tariff policies, i.e., policies that require detailed and frequent consumption measurements, in order to determine the bill. Meters do not send consumption measurements to the service provider. Instead, the computation of the bill is done locally and only the amount to be paid is revealed to the service provider. We improve existing work in two ways. First, we generalize the security model to consider multiple meters and multiple users. Second, we propose a privacy-preserving billing protocol for our model that improves efficiency for policies described by splines.

Computations on signed data. We have shown how the service provider can perform computations on unencrypted signed data.

Data anonymization and sharing. We have proposed an (un)linkable pseudonym system to allow exchange of user data between databases. A converter serves as central hub to ensure controllability. The converter establishes individual pseudonyms for each server derived from a unique main identifier that every user has, but without learning the derived pseudonyms. The only information the converter still learns is that a server S_A wants to access data from a server S_B, which is the right amount of information to balance control and privacy.

In Task 24.4, we have designed several privacy-friendly revocation mechanisms. First, we have proposed a privacy preserving revocation mechanism for privacy-enhancing attribute-based credentials that allows you to efficiently handle multiple revocation lists. Second, we have studied a primitive that is widely used for revocation purposes, i.e., cryptographic accumulators. Third,

we have shown how using epochs can help to make revocation practical while still retaining reasonable strong privacy guarantees. Finally, we have explained the concept of revocable privacy.

In Task 24.5, we have designed a two-factor user-authentication scheme for usable server-based eID and e-signature solutions. Current server-based eID and e-signature solutions typically rely on one-time passwords delivered to the user via short message service (SMS). This raises several issues in practice, as the use of SMS technology can be cost-effective insecure. To address these issues, we have proposed an alternative two-factor user-authentication scheme following a challenge-response approach. The feasibility and applicability of the proposed user-authentication scheme has been evaluated by means of two concrete implementations. This way, we have shown that the proposed authentication scheme and its implementations improve both the cost effectiveness and the security of server-based eID and e-signature solutions. Additionally, on a different line of work, we have studied how users choose passwords under consideration of different human dimensions, and, more specifically, when they are cognitively depleted.

10 Abstracts of Research Papers in D24.4

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

- Omar Almousa, Sebastian Mödersheim, Paolo Modesti, and Luca Viganò: *Typing and Compositionality for Security Protocols: A Generalization to the Geometric Fragment*, in *ESORICS 2015* [26].

  We integrate, and improve upon, prior relative soundness results of two kinds. The first kind are typing results showing that any security protocol that fulfils a number of sufficient conditions has an attack if it has a well-typed attack. The second kind considers the parallel composition of protocols, showing that when running two protocols in parallel allows for an attack, then at least one of the protocols has an attack in isolation. The most important generalization over previous work is the support for all security properties of the geometric fragment.

- Paolo Modesti: *AnBx: Automatic Generation and Verification of Security Protocols Implementations*, in *FPS 2015*, to appear [269].

  The AnBx compiler is a tool for automatic generation of Java implementations of security protocols specified in a simple and abstract model that can be formally verified. In our model-driven development approach, protocols are described in AnBx , an extension of the Alice & Bob notation; along with the synthesis of consistency checks, the tool analyses the security goals and produces annotations that allow the verification of the generated implementation with ProVerif.

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

- Sebastian Mödersheim, Omar Almousa, Bud Bruegger, Max Tuengerthal: "A Formal Verification of the FutureID architecture". DTU technical report, 2015 [263].

  Abstract: FutureID implements an identity management infrastructure able to support the European single market. At this scale, requirements for security, privacy, and accountability are ever more stringent. As a major contribution to reaching these requirements, we formalize and formally verify the FutureID architecture on an abstract level, in particular for the three goals. This has even lead to the development of new protocols and possibilities for future FutureID.

- Jan Camenisch, Maria Dubovitskaya, Robert R. Enderlein, Anja Lehmann, Gregory Neven, Christian Paquin, and Franz-Stefan Preiss: "Concepts and Languages for Privacy Preserving Attribute-Based Authentication", in *Journal of Information Security and Applications* [90].

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), provide a language framework in XML schema, and give a formal semantics to describe the effects of the transactions in a privacy-friendly authentication system using Privacy-ABCs. 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 algorithms either.

- Jan Camenisch, Sebastian Mödersheim, Gregory Neven, Franz-Stefan Preiss, and Alfredo Rial: "A Prolog Program for Matching Attribute-Based Credentials to Access Control Policies", in *IBM Research Report RZ3890* [119].

  In an attribute-based credential system, users employ credentials issued by credentials issuers to compute presentation tokens that prove to service providers that the user's credentials fulfill the access control policies to access services. The number of user credentials and the number of ways a policy can be satisfied can be large. Therefore, a user has to choose which subset of her credentials she wishes to employ to compute a presentation token. This choice has both efficiency and privacy implications. We present a Prolog program that lists all the credentials subsets that can be used to fulfill a given policy. In our program, credentials are represented by facts and policies by rules. By querying a rule, the Prolog engine lists all the combinations of facts that satisfy the rule. Therefore, we remark the simplicity of our approach, which simply requires representing credentials and policies in Prolog and avoids the need of implementing credential-policy matching or exhaustive search algorithms. Furthermore, our program is also useful on the verifier side. By using facts to represent the credential information disclosed by a user's presentation tokens, when the user wishes to access a new service, the service provider can verify whether the credential information already disclosed fulfills the policy for that service. Our Prolog program implements a variety of features of an attribute-based credential system: pseudonyms, key binding, different restrictions for attributes values, issuer-driven and verifier-driven revocation, and inspection. Our program can easily be extended to implement more features.

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

- Thomas Groß: "Signatures and Efficient Proofs on Committed Graphs and NP-Statements", in *FC 2015* [197]

  Digital signature schemes are a foundational 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 (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 establish 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 Σ-proofs and the strong RSA assumption. 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.

- Jan Camenisch, Robert Enderlein, and Gregory Neven: "Two-Server Password Authenticated Secret Sharing UC-Secure Against Transient Corruptions", in *PKC 2015* [97]

  Protecting user data entails providing authenticated users access to their data. The most prevalent and probably also the most feasible approach to the latter is by username and password. With password breaches through server compromise now reaching billions of affected passwords, distributing the password files and user data over multiple servers is not just a good idea, it is a dearly needed solution to a topical problem. Threshold password-authenticated secret sharing (TPASS) protocols enable users to share secret data among a set of servers so that they can later recover that data using a single password. No coalition of servers up to a certain threshold can learn anything about the data or perform an offline dictionary attack on the password. Several TPASS protocols have appeared in the literature and one is even available commercially. Although designed to tolerate corrupted servers, unfortunately none of these protocols provide details let alone security proofs about the steps that need to be taken when a compromise actually occurs and how to proceed. Indeed, they consider static corruptions only which for instance does not model real world attacks by hackers. We provide the first TPASS protocol that is provably secure against adaptive server corruptions. Moreover, our protocol contains an efficient recovery procedure allowing one to re-initialize servers to recover from corruption. We prove our protocol secure in the universal composability model where servers can be corrupted adaptively at any time; the users' passwords and secrets remain safe as long as both servers are not corrupted at the same time. Our protocol does not require random oracles but does assume that servers have certified public keys.

- Jan Camenisch, Anja Lehmann and Gregory Neven: "Optimal Distributed Password Verification", in *ACM CCS 2015* [112]

  We present a highly efficient cryptographic protocol to protect user passwords against server compromise by distributing the capability to verify passwords over multiple servers.

Password verification is a single-round protocol and requires from each server only one exponentiation in a prime-order group. In spite of its simplicity, our scheme boasts security against dynamic and transient corruptions, meaning that servers can be corrupted at any time and can recover from corruption by going through a non-interactive key refresh procedure. The users' passwords remain secure against offline dictionary attacks as long as not all servers are corrupted within the same time period between refreshes. The only currently known scheme to achieve such strong security guarantees incurs the considerable cost of several hundred exponentiations per server. We prove our scheme secure in the universal composability model, which is well-known to offer important benefits for password-based primitives, under the gap one-more Diffie-Hellman assumption in the random-oracle model. Server initialization and refresh must take place in a trusted execution environment. Initialization additionally requires a secure message to each server, but the refresh procedure is non-interactive. We show that these requirements are easily met in practice by providing an example deployment architecture.

- Jan Camenisch, Stephan Krenn, Anja Lehmann, Gert Læssøe Mikkelsen, Gregory Neven and Michael Østergaard Pedersen: "Formal Treatment of Privacy-Enhancing Credential Systems", in *SAC 2015* [107]

  Privacy-enhancing attribute-based credentials (PABCs) are the core ingredients to privacy-friendly authentication systems. They allow users to obtain credentials on attributes and prove possession of these credentials in an unlinkable fashion while revealing only a subset of the attributes. In practice, PABCs typically need additional features like revocation, pseudonyms as privacy-friendly user public keys, or advanced issuance where attributes can be "blindly" carried over into new credentials. For many such features, provably secure solutions exist in isolation, but it is unclear how to securely combined them into a full-fledged PABC system, or even which properties such a system should fulfill.

  We provide a formal treatment of PABCs supporting a variety of features by defining their syntax and security properties, resulting in the most comprehensive definitional framework for PABCs so far. Unlike previous efforts, our definitions are not targeted at one specific use-case; rather, we try to capture generic properties that can be useful in a variety of scenarios. We believe that our definitions can also be used as a starting point for diverse application-dependent extensions and variations of PABCs. We present and prove secure a generic and modular construction of a PABC system from simpler building blocks, allowing for a "plug-and-play" composition based on different instantiations of the building blocks. Finally, we give secure instantiations for each of the building blocks.

- Jan Camenisch, Maria Dubovitskaya, Kristiyan Haralambiev and Markulf Kohlweiss: "Unlinkable Redactable Signatures and Their Applications to Anonymous Credentials", in *ASIACRYPT 2015* [92]

  We review the design principles of anonymous credentials and similar privacy-enhancing protocols and conclude that theoretical cryptographic advances have not been sufficiently applied to the construction of practical protocols. Arguably, all schemes suited for real-world-use do not support straight-line extraction and most of them still rely on random oracles in their security arguments (because they use the Fiat-Shamir heuristic to obtain non-interactive proofs). To address this gap we propose a new kind of signature

scheme, an *unlinkable redactable block-signature* (URS) scheme as a new building block for privacy-enhancing protocols. We give both property-based security definitions and a UC functionality for URS and validate the former by showing that they imply the latter. We then employ structure preserving signatures and vector commitments to construct a concrete URS scheme and prove that it meets our property-based security definitions.

Finally, we show how our new signature scheme can be used to construct anonymous credentials and create the first efficient UC-secure credential scheme for which both the size of a credential and its presentation proof are independent of the number of attributes issued in a credential. Moreover, our new credential scheme is secure under DH-like cryptographic assumptions, i.e., it does not rely on random oracles. Our definitional approach may be of independent interest beyond redactable signatures as a framework for UC-secure malleable signatures and we expect our URS scheme to be useful for the construction of many privacy-preserving protocols.

- Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig: "Practical Round-Optimal Blind Signatures in the Standard Model", in *CRYPTO 2015* [183]

  Round-optimal blind signatures are notoriously hard to construct in the standard model—especially in the malicious-signer model, where blindness has to hold even under adversarially chosen keys. This practical experience is substantiated by several impossibility results. So far, the only construction that can be termed theoretically efficient, by Garg and Gupta (Eurocrypt'14), requires complexity leveraging, which induces an exponential security loss.

  In this paper we present the first construction of practically efficient round-optimal blind signatures in the standard model. It is conceptually simple and builds on the recent structure preserving signatures on equivalence classes (SPS-EQ) from Asiacrypt'14. While the traditional notion of blindness follows from the security of the SPS-EQ and standard assumptions, we prove blindness under adversarially chosen keys under an interactive variant of DDH. In contrast to previous constructions, we neither require non-uniform assumptions nor complexity leveraging.

  We also show how to extend our construction to partially blind signatures and to blind signatures on message vectors, which directly yields the first standard-model construction of one-show anonymous credentials à la "anonymous credentials light" (CCS'13).

  Applying Fischlin and Schröder's impossibility result for blind signatures (Eurocrypt'10) to our construction gives us new insights on standard-model constructions of SPS-EQ. Furthermore, we provide the first standard-model SPS-EQ construction and show how SPS-EQ schemes imply conventional structure-preserving signatures.

- David Derler, Christian Hanser, and Daniel Slamanig: "Blank Digital Signatures: Optimization and Practical Experiences", in *Privacy and Identity Management for the Future Internet in the Age of Globalisation* [163]

  Blank Digital Signatures (BDS) [206] enable an originator to delegate the signing rights for a template, containing fixed and exchangeable elements, to a proxy. The proxy is then able to choose one of the predefined values for each exchangeable element and issue a signature for such an instantiation of the template on behalf of the originator. In this paper, we

propose optimizations for the BDS scheme from [206] and present a library, integrating this optimized version within the Java Cryptography Architecture and the keying material into X.509 certificates. To illustrate the flexibility of the proposed library, we introduce two proof-of-concept implementations building up on XML and PDF, respectively. Finally, we give a detailed insight in the performance of the protocol and our implementation.

- Bernd Zwattendorfer, and Daniel Slamanig: "Design strategies for a privacy-friendly Austrian eID system in the public cloud", in *Computers & Security* [352]

  Abstract Secure identification and authentication are essential processes in sensitive areas of application such as e-Government or e-Health. In Austria, the official eID is the so called the Austrian citizen card and a means of choice for secure citizen identification and authentication. To facilitate the adoption of citizen card authentication at service providers within the Austrian e-Government strategy, the open source module MOA-ID has been developed. It acts as identity provider for different service providers and manages identification and authentication based on the Austrian citizen card. Currently, MOA-ID is deployed locally in every service provider's domain and is assumed to be fully trusted. With the increasing use of eIDs, however, a move into a public cloud might be advantageous due to benefits provided cloud computing, e.g., cost savings or scalability. Nevertheless, the move of a trusted service into the public cloud brings up new obstacles, in particular with respect to security and privacy. Therefore, in this paper we introduce and evaluate three different approaches on how the Austrian eID system based on MOA-ID could be securely moved into the cloud without violating any privacy or data protection aspects. To achieve this, we rely on various cryptographic methods and focus on minimum changes of the current identification and authentication process flow. Based on an evaluation of these three different approaches, we propose a model which can be generically used for eID identification and authentication in privacy-invasive environments such as the public cloud.

- Jan Camenisch, Diego A. Ortiz-Yepes and Franz-Stefan Preiss: "Strengthening Authentication with Privacy-Preserving Location Verification of Mobile Phones", in *WPES 2015* [123]

  Mobile devices are increasingly used in security-sensitive contexts such as physical access control and authorization of payment transactions. In this paper we contribute a mechanism to verify whether a mobile device currently resides within a geographical area at a given time, thus enabling the use of the location as an additional authentication factor. Trustworthiness, privacy, and practicability are central to our mechanism. In particular, to provide trustworthy location information, our mechanism uses the location of the phone *as detected by the Mobile Network Operator* instead of relying on the location detected by the phone itself, which can be manipulated. We have followed a privacy-by-design approach to ensure that sensitive information, e.g., location and subscriber data, are only revealed to parties with a need to know. Privacy safeguards are realized using anonymous credentials, an established privacy-enhancing technology. Finally, our mechanism is practical and has little requirements on the mobile phone beyond the ability to run computations on anonymous credentials, as well as Internet and mobile network connectivity. These requirements are fulfilled by most smartphones in the market.

- Alfredo Rial: "Blind Attribute-Based Encryption and Oblivious Transfer with Fine-Grained Access Control", in *Designs, Codes and Cryptography* [304]

  We propose two constructions of oblivious transfer with access control (OTAC), i.e., oblivious transfer schemes in which a receiver can obtain a message only if her attributes, which are certified by a credential issuer, satisfy the access control policy of that message. The receiver remains anonymous towards the sender and the receiver's attributes are not disclosed to the sender. Our constructions are based on any ciphertext policy attribute based encryption (CPABE) scheme that fulfills the committing and key separation properties, which we define. We also provide a committing CPABE with key separation scheme that supports any policy described by a monotone access structure, which, in comparison to previous work, allows our OTAC construction to support efficiently a wider variety of access control policies. In our constructions, a receiver obtains from the sender a CPABE secret key for her attributes by using a blind key extraction with access control protocol. We provide a blind key extraction with access control protocol for any committing CPABE with key separation scheme. Previous work only provided ad-hoc constructions of blind key extraction protocols. Our generic protocol works in a hybrid model that employs novel ideal functionalities for oblivious transfer and for anonymous attribute authentication. We propose constructions that realize those novel ideal functionalities and analyze the overall efficiency of our OTAC constructions.

- Alfredo Rial, George Danezis and Markulf Kohlweiss: "Privacy-Preserving Smart Metering Revisited" [306]

  Privacy-preserving billing protocols are useful in settings where a meter measures user consumption of some service, such as smart metering of utility consumption, pay-as-you-drive insurance and electronic toll collection. In such settings, service providers apply fine-grained tariff policies that require meters to provide a detailed account of user consumption. The protocols allow the user to pay to the service provider without revealing the user's consumption measurements. Our contribution is twofold. First, we propose a general model where a meter can output meter readings to multiple users, and where a user receives meter readings from multiple meters. Unlike previous schemes, our model accommodates a wider variety of smart metering applications. Second, we describe a protocol based on polynomial commitments that improves the efficiency of previous protocols for tariff policies that employ splines to compute the price due.

- Jae Hyun Ahn, Dan Boneh, Jan Camenisch, Susan Hohenberger, Abhi Shelat, and Brent Waters: "Computing on Authenticated Data", in *Journal of Cryptology* [23]

  In tandem with recent progress on computing on encrypted data via fully homomorphic encryption, we present a framework for computing on *authenticated* data via the notion of slightly homomorphic signatures, or P-homomorphic signatures. With such signatures, it is possible for a third party to *derive* a signature on the object m' from a signature of m as long as $P(m, m') = 1$ for some predicate P which captures the "authenticatable relationship" between m' and m. Moreover, a derived signature on m' reveals *no extra information* about the parent m.

  Our definition is carefully formulated to provide one unified framework for a variety of

distinct concepts in this area, including arithmetic, homomorphic, quotable, redactable, transitive signatures and more. It includes being unable to distinguish a derived signature from a fresh one *even when given the original signature*. The inability to link derived signatures to their original sources prevents some practical privacy and linking attacks, which is a challenge not satisfied by most prior works.

Under this strong definition, we then provide generic constructions for all univariate and closed predicates, and specific efficient constructions for a broad class of natural predicates such as quoting, subsets, weighted sums, averages, and Fourier transforms. To our knowledge, these are the first efficient constructions for these predicates (excluding subsets) that provably satisfy this strong security notion.

- Jan Camenisch and Anja Lehmann: "(Un)linkable Pseudonyms for Governmental Databases", in *ACM CCS 2015* [110]

  When data maintained in a decentralized fashion needs to be synchronized or exchanged between different databases, related data sets usually get associated with a unique identifier. While this approach facilitates cross-domain data exchange, it also comes with inherent drawbacks in terms of controllability. As data records can easily be linked, no central authority can limit or control the information flow. Worse, when records contain sensitive personal data, as is for instance the case in national social security systems, such linkability poses a massive security and privacy threat. An alternative approach is to use domain-specific pseudonyms, where only a central authority knows the cross-domain relation between the pseudonyms. However, current solutions require the central authority to be a fully trusted party, as otherwise it can provide false conversions and exploit the data it learns from the requests. We propose an (un)linkable pseudonym system that overcomes those limitations, and enables controlled yet privacy-friendly exchange of distributed data. We prove our protocol secure in the UC framework and provide an efficient instantiation based on discrete-logarithm related assumptions.

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

- Jan Camenisch, Maria Dubovitskaya, and Alfredo Rial: "UC Commitments, Revocation, and Attribute Tokens for Privacy-Preserving Protocol Design" [96]

  Attribute-based credentials can be revoked in some context, while still being valid for other purposes. We propose two accumulator-based revocation schemes from vector commitments that allow one to accumulate several revocation lists into a single commitment value, and allow a user to have only one witness for those several revocation lists. One of our schemes, unlike standard accumulators, hides the revocation status of a user from other users and verifiers.

  To make our revocation building blocks employable by other protocols, we define the ideal revocation functionalities for both hiding and non-hiding settings in the universal composability framework and prove that our schemes realize them. We also provide a new commitment functionality that outputs cryptographic commitments and, therefore, can be employed in a hybrid protocol together with non-interactive proof-of-knowledge

functionalities that prove different statements about the committed value. We believe that our work on the composable trapdoor commitment schemes is of independent interest and can be used outside of the revocation context. Finally, we demonstrate the feasibility of our approach by constructing universally composable anonymous attribute tokens that support efficient revocation in multiple contexts.

- David Derler, Christian Hanser, and Daniel Slamanig: "Revisiting Cryptographic Accumulators, Additional Properties and Relations to other Primitives", in *Topics in Cryptology - CT-RSA 2015* [165]

  Cryptographic accumulators allow to accumulate a finite set of values into a single succinct accumulator. For every accumulated value, one can efficiently compute a witness, which certifies its membership in the accumulator. However, it is computationally infeasible to find a witness for any non-accumulated value. Since their introduction, various accumulator schemes for numerous practical applications and with different features have been proposed. Unfortunately, to date there is no unifying model capturing all existing features. Such a model can turn out to be valuable as it allows to use accumulators in a black-box fashion.

  To this end, we propose a unified formal model for (randomized) cryptographic accumulators which covers static and dynamic accumulators, their universal features and includes the notions of undeniability and indistinguishability. Additionally, we provide an exhaustive classification of all existing schemes. In doing so, it turns out that most accumulators are distinguishable. Fortunately, a simple, light-weight generic transformation allows to make many existing dynamic accumulator schemes indistinguishable. As this transformation, however, comes at the cost of reduced collision freeness, we additionally propose the first indistinguishable scheme that does not suffer from this shortcoming. Finally, we employ our unified model for presenting a black-box construction of commitments from indistinguishable accumulators as well as a black-box construction of indistinguishable, undeniable universal accumulators from zero-knowledge sets. Latter yields the first universal accumulator construction that provides indistinguishability.

- Wouter Lueks, Gergely Alpár, Jaap-Henk Hoepman, and Pim Vullers: "Fast Revocation of Attribute-Based Credentials for Both Users and Verifiers" in *ICT Systems Security and Privacy Protection - 30th IFIP TC 11 International Conference, SEC 2015, Hamburg, Germany, May 26-28, 2015* [246]

  Attribute-based credentials allow a user to prove properties about herself anonymously. Revoking such credentials, which requires singling them out, is hard because it is at odds with anonymity. All revocation schemes proposed to date either sacrifice anonymity altogether, require the parties to be online, or put high load on the user or the verifier. As a result, these schemes are either too complicated for low-powered devices like smart cards or they do not scale. We propose a new revocation scheme that has very low computational cost for users and verifiers, and that does not require users to process updates. We trade only a limited amount of anonymity—in exceptional cases, uses of credentials are linkable—to make the first practical revocation scheme that is efficient at large scales and fast enough for smart cards.

- Wouter Lueks, Maarten H. Everts and Jaap-Henk Hoepman: "Revocable Privacy: Principles, Use Cases, and Technologies" in *Annual Privacy Forum (APF 2015), Luxembourg, October 7-8 2015. (to appear)* [247].

  Security and privacy often seem to be at odds with one another. In this paper, we revisit the design principle of revocable privacy which guides the creation of systems that offer anonymity for people who do not violate a predefined rule, but can still have consequences for people who do violate the rule. We first improve the definition of revocable privacy by considering different types of sensors for users' actions and different types of consequences of violating the rules (for example blocking). Second, we explore some use cases that can benefit from a revocable privacy approach. For each of these, we derive the underlying abstract rule that users should follow. Finally, we describe existing techniques that can implement some of these abstract rules. These descriptions not only illustrate what can already be accomplished using revocable privacy, they also reveal directions for future research.

10.5 Methods for Usable Privacy (Task 24.5)

- Christof Rath, Simon Roth, Harald Bratko, and Thomas Zefferer: "Encryption-based Second Authentication Factor Solutions for Qualified Server-side Signature Creation", in *EGOVIS 2015* [302]

  Electronic identity (eID) and electronic signature (e-signature) are key concepts of transactional e-government solutions. Especially in Europe, server-based eID and e-signature solutions have recently gained popularity, as they provide enhanced usability while still complying with strict security requirements. To implement obligatory two-factor user-authentication schemes, current server-based eID and e-signature solutions typically rely on one-time passwords delivered to the user via short message service (SMS). This raises several issues in practice, as the use of SMS technology can be cost-effective insecure. To address these issues, we propose an alternative two-factor user-authentication scheme following a challenge-response approach. The feasibility and applicability of the proposed user-authentication scheme is evaluated by means of two concrete implementations. This way, we show that the proposed authentication scheme and its implementations improve both the cost effectiveness and the security of server-based eID and e-signature solutions.

11 List of Research Papers in WP24

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

- Sebastian Mödersheim and Georgios Katsoris. "A Sound Abstraction of the Parsing Problem" in *Computer Security Foundations 2014* [266].
- Sebastian Mödersheim and Luca Viganò: "Sufficient Conditions for Vertical Composition of Security Protocols" in *AsiaCCS 2014* [267].
- Paolo Modesti, "Efficient Java Code Generation of Security Protocols specified in AnB/AnBx" in *STM 2014* [270].
- Omar Almousa, Sebastian Mödersheim, Paolo Modesti, and Luca Viganò: "Typing and Compositionality for Security Protocols: A Generalization to the Geometric Fragment" in *ESORICS 2015* [26].
- Paolo Modesti, "AnBx: Automatic Generation and Verification of Security Protocols Implementations" in *FPS 2015* [269].

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

- Sebastian Mödersheim, Thomas Groß, and Luca Viganò: "Defining Privacy is Supposed to be Easy", LPAR 2013 [265].
- Sebastian Mödersheim, Omar Almousa, Bud Bruegger, Max Tuengerthal: "A Formal Verification of the FutureID architecture". DTU technical report, 2015 [263].
- Jan Camenisch, Maria Dubovitskaya, Robert R. Enderlein, Anja Lehmann, Gregory Neven, Christian Paquin, and Franz-Stefan Preiss: "Concepts and Languages for Privacy Preserving Attribute-Based Authentication", in *Journal of Information Security and Applications* [90].
- Jan Camenisch, Sebastian Mödersheim, Gregory Neven, Franz-Stefan Preiss, and Alfredo Rial: "A Prolog Program for Matching Attribute-Based Credentials to Access Control Policies", in *IBM Research Report RZ3890* [119].

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

- Jan Camenisch and Anja Lehmann: "(Un)linkable Pseudonyms for Governmental Databases", in *ACM CCS 2015* [110]
- Jan Camenisch, Anja Lehmann and Gregory Neven: "Optimal Distributed Password Verification", in *ACM CCS 2015* [112]

- Jan Camenisch, Stephan Krenn, Anja Lehmann, Gert Læssøe Mikkelsen, Gregory Neven and Michael Østergaard Pedersen: "Formal Treatment of Privacy-Enhancing Credential Systems", in *SAC 2015* [107]
- Jan Camenisch, Diego A. Ortiz-Yepes and Franz-Stefan Preiss: "Strengthening Authentication with Privacy-Preserving Location Verification of Mobile Phones", in *WPES 2015* [123]
- Jan Camenisch, Maria Dubovitskaya, Kristiyan Haralambiev and Markulf Kohlweiss: "Unlinkable Redactable Signatures and Their Applications to Anonymous Credentials", in *WPES 2015* [92]
- Alfredo Rial: "Blind Attribute-Based Encryption and Oblivious Transfer with Fine-Grained Access Control", in *Designs, Codes and Cryptography* [304]
- Alfredo Rial, George Danezis and Markulf Kohlweiss: "Privacy-Preserving Smart Metering Revisited" [306]
- Christian Hanser and Daniel Slamanig, "Blank Digital Signatures," in *Proc. of 8th ACM Symposium on Information, Computer and Communications Security (ASIACCS 2013)* [206].
- 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)* [207].
- 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)* [350].
- 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)* [351].
- Markulf Kohlweiss and Alfredo Rial, "Optimally Private Access Control," in *Workshop on Privacy in the Electronic Society 2013 (WPES 2013)* [233].
- Jan Camenisch and Robert Enderlein and Victor Shoup, "Practical and Employable Protocols for UC-Secure Circuit Evaluation over Z_n," in *ESORICS 2013* [98].
- 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* [12].
- Changyu Dong and Liqun Chen and Jan Camenisch and Giovanni Russello, "Fair Private Set Intersection with a Semi-trusted Arbiter," in *DBSec 2013* [168].
- Fabrice Benhamouda and Jan Camenisch and Stephan Krenn and Vadim Lyubashevsky and Gregory Neven, "Better Zero-Knowledge Proofs for Lattice Encryption and Their Application to Group Signatures," in *ASIACRYPT 2014* [56].

- Jan Camenisch and Günter Karjoth and Gregory Neven and Franz-Stefan Preiss, "Anonymously sharing Flickr pictures with facebook friends," in *Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society (WPES 2013)* [102].
- Jan Camenisch and Anja Lehmann and Gregory Neven and Alfredo Rial, "Privacy Preserving Auditing for Attribute-Based Credentials," in *Computer Security - ESORICS 2014* [113].
- Daniel Slamanig and Klaus Stranacher and Bernd Zwattendorfer, "User-Centric Identity as a Service-Architecture for eIDs with Selective Attribute Disclosure" in *ACM Symposium on Access Control Models and Technologies - SACMAT 2014* [321].
- Christian Hanser and Daniel Slamanig, "Structure-Preserving Signatures on Equivalence Classes and their Application to Anonymous Credentials" in *ASIACRYPT 2014* [208].
- David Derler and Christian Hanser and Daniel Slamanig, "Privacy-Enhancing Proxy Signatures from Non-Interactive Anonymous Credentials" in *Data and Applications Security and Privac - DBSec 2014* [164]
- Thomas Groß, "Efficient Certification and Zero-Knowledge Proofs of Knowledge on Infrastructure Topology Graphs" in *ACM CCS Cloud Security Workshop (CCSW) 2014* [196]
- Jae Hyun Ahn, Dan Boneh, Jan Camenisch, Susan Hohenberger, Abhi Shelat, and Brent Waters: "Computing on Authenticated Data", in *Journal of Cryptology* [23].
- Thomas Groß, "Signatures and Efficient Proofs on Committed Graphs and NP-Statements" in *FC 2015* [197].
- Jan Camenisch, Robert Enderlein, and Gregory Neven: "Two-Server Password Authenticated Secret Sharing UC-Secure Against Transient Corruptions", in *PKC 2015* [97].
- Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig: "Practical Round-Optimal Blind Signatures in the Standard Model", in *CRYPTO 2015* [183].
- David Derler, Christian Hanser, and Daniel Slamanig: "Blank Digital Signatures: Optimization and Practical Experiences", in *Privacy and Identity Management for the Future Internet in the Age of Globalisation* [163].
- Bernd Zwattendorfer, and Daniel Slamanig: "Design strategies for a privacy-friendly Austrian eID system in the public cloud", in *Computers & Security* [352].

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

- Daniel Slamanig and Raphael Spreitzer and Thomas Unterluggauer, "Adding Controllable Linkability to Pairing-Based Group Signatures For Free," in *International Conference on Information Security - ISC 2014* [320].
- Jan Camenisch, Maria Dubovitskaya, and Alfredo Rial: "UC Commitments, Revocation, and Attribute Tokens for Privacy-Preserving Protocol Design" [96].

- David Derler, Christian Hanser, and Daniel Slamanig: "Revisiting Cryptographic Accumulators, Additional Properties and Relations to other Primitives", in *Topics in Cryptology - CT-RSA 2015* [165].

- Wouter Lueks, Gergely Alpár, Jaap-Henk Hoepman, and Pim Vullers: "Fast Revocation of Attribute-Based Credentials for Both Users and Verifiers" in *ICT Systems Security and Privacy Protection - 30th IFIP TC 11 International Conference, SEC 2015, Hamburg, Germany, May 26-28, 2015* [246].

- Wouter Lueks, Maarten H. Everts and Jaap-Henk Hoepman: "Revocable Privacy: Principles, Use Cases, and Technologies" in *Annual Privacy Forum (APF 2015), Luxembourg, October 7-8 2015. (to appear)* [247].

11.5 Methods for Usable Privacy (Task 24.5)

- Kovila P.L. Coopamootoo and Thomas Groß, "Mental Models for Usable Privacy: A Position Paper", in *Proceedings of the Human Factors in Security and Privacy Conference, HCI International Conference (HAS2014)* [151].

- Kovila P.L. Coopamootoo and Thomas Groß, "Poster: Preliminary Investigation of Cognitive Effort in Privacy Decision-Making: Personal Information vs. 3 x 4", at the *Symposium on Usable Security and Privacy (SOUPS2014)* [153].

- Kovila P.L. Coopamootoo and Thomas Groß, "Mental Models: An Approach to Identify Privacy Concern and Behavior", at the *Privacy Personas and Segmentation Workshop (SOUPS2014)* [149].

- Kovila P.L. Coopamootoo and Thomas Groß, "Mental Models of Online Privacy: Structural Properties with Cognitive Maps",WIPS in *Proceedings of the 28th British Human Computer Interaction Conference (BCS HCI2014)* [152].

- Kovila P.L. Coopamootoo and Thomas Groß, "Cognitive Effort in Privacy Decision-Making vs. 3 x 4: Evaluation of a Pilot Experiment Design", to appear at *Learning from Authoritative Security Experiments Workshop (LASER2014)* [150].

- Christof Rath, Simon Roth, Harald Bratko, and Thomas Zefferer: "Encryption-based Second Authentication Factor Solutions for Qualified Server-side Signature Creation", in *EGOVIS 2015* [302].

List of References

[1] Austrian citizen card. Technical report. URL: http://www.a-sit.at/de/dokumente_publikationen/flyer/buergerkarte_en.php.

[2] Belgian crossroads bank for social security. Technical report. URL: http://www.ksz.fgov.be/.

[3] Directive 2009/136/ec. Official Journal of the European Union (2009).

[4] Irma – i reveal my attributes. Research Project.

[5] Proof of concept on integrating german identity scheme with u-prove technology.

[6] Regulation (ec) no 45/2001. Official Journal of the European Union (2001).

[7] Harald Aamot, Christian D Kohl, Daniela Richter, and Petra Knaup-Gregori. Pseudonymization of patient identifiers for translational research. *BMC medical informatics and decision making*, 13(1):75, 2013.

[8] 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.

[9] Michel Abdalla, Chanathip Namprempre, and Gregory Neven. On the (im)possibility of blind message authentication codes. In David Pointcheval, editor, *Topics in Cryptology – CT-RSA 2006*, volume 3860 of *Lecture Notes in Computer Science*, pages 262–279. Springer, February 2006.

[10] Masayuki Abe. A Secure Three-Move Blind Signature Scheme for Polynomially Many Signatures. In Pfitzmann [295], pages 136–151.

[11] Masayuki Abe, editor. *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*. Springer, 2010.

[12] Masayuki Abe, Jan Camenisch, Maria Dubovitskaya, and Ryo Nishimaki. Universally composable adaptive oblivious transfer (with access control) from standard assumptions. In *Proceedings of the 2013 ACM workshop on Digital identity management*, pages 1–12. ACM, 2013.

[13] Masayuki Abe, Melissa Chase, Bernardo David, Markulf Kohlweiss, Ryo Nishimaki, and Miyako Ohkubo. Constant-size structure-preserving signatures: Generic constructions and simple assumptions. In *Advances in Cryptology – ASIACRYPT 2012*, Lecture Notes in Computer Science, pages 4–24. Springer, December 2012. doi:10.1007/978-3-642-34961-4_3.

[14] Masayuki Abe, Georg Fuchsbauer, Jens Groth, Kristiyan Haralambiev, and Miyako Ohkubo. Structure-preserving signatures and commitments to group elements. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 209–236. Springer, August 2010.

[15] Masayuki Abe, Jens Groth, Kristiyan Haralambiev, and Miyako Ohkubo. Optimal structure-preserving signatures in asymmetric bilinear groups. In Phillip Rogaway, editor, *Advances in Cryptology – CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 649–666. Springer, August 2011.

[16] Masayuki Abe, Jens Groth, Miyako Ohkubo, and Mehdi Tibouchi. Unified, minimal and selectively randomizable structure-preserving signatures. In *TCC 2014: 11th Theory of Cryptography Conference*, Lecture Notes in Computer Science, pages 688–712. Springer, 2014. doi:10.1007/978-3-642-54242-8_29.

[17] Masayuki Abe, Kristiyan Haralambiev, and Miyako Ohkubo. Group to group commitments do not shrink. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 301–317. Springer, April 2012.

[18] Masayuki Abe and Miyako Ohkubo. A framework for universally composable non-committing blind signatures. In Mitsuru Matsui, editor, *Advances in Cryptology – ASIACRYPT 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 435–450. Springer, December 2009.

[19] Masayuki Abe and Tatsuaki Okamoto. Provably secure partially blind signatures. In Mihir Bellare, editor, *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 271–286. Springer, August 2000.

[20] Tolga Acar and Lan Nguyen. Revocation for delegatable anonymous credentials. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011: 14th* , volume 6571 of *Lecture Notes in Computer Science*, pages 423–440. Springer, March 2011.

[21] 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.

[22] 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.

[23] Jae Hyun Ahn, Dan Boneh, Jan Camenisch, Susan Hohenberger, Abhi Shelat, and Brent Waters. Computing on authenticated data. *Journal of Cryptology*, pages 1–45.

[24] Jae Hyun Ahn, Dan Boneh, Jan Camenisch, Susan Hohenberger, Abhi Shelat, and Brent Waters. Computing on authenticated data. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 1–20. Springer, March 2012.

[25] Omar Almousa, Sebastian Mödersheim, and Luca Viganò. Alice and bob: Reconciling formal models and implementation. In Chiara Bodei, Gian-Luigi Ferrari, and Corrado Priami, editors, *Programming Languages with Applications to Biology and Security*, volume 9465 of *Lecture Notes in Computer Science*, pages 66–85. Springer International Publishing, 2015. URL: http://dx.doi.org/10.1007/978-3-319-25527-9_7, doi:10.1007/978-3-319-25527-9_7.

[26] Omar Almousa, Sebastian Mödersheim, Paolo Modesti, and Luca Viganò. Typing and compositionality for security protocols: A generalization to the geometric fragment. In *Computer Security - ESORICS 2015 - 20th European Symposium on Research in Computer Security, Vienna, Austria, September 21-25, 2015, Proceedings, Part II*, pages 209–229, 2015. URL: http://dx.doi.org/10.1007/978-3-319-24177-7_11, doi:10.1007/978-3-319-24177-7_11.

[27] Omar Almousa, Sebastian Mödersheim, Paolo Modesti, and Luca Viganò. Typing and Compositionality for Security Protocols: A Generalization to the Geometric Fragment (Extended Version). Technical Report DTU Compute TR-2015-03, Technical University of Danmark, 2015. Available at http://compute.dtu.dk/~samo.

[28] Ross Anderson and Shailendra Fuloria. On the security economics of electricity metering. In *WEIS*, 2010.

[29] 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.

[30] Myrto Arapinis and Marie Duflot. Bounding messages for free in security protocols - extension to various security properties. *Inf. Comput.*, 239:182–215, 2014.

[31] A. Armando and L. Compagna. SATMC: A SAT-based Model Checker for Security Protocols. *Logics in Artificial Intelligence*, pages 730–733, 2004.

[32] 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.

[33] Giuseppe Ateniese, Dawn Xiaodong Song, and Gene Tsudik. Quasi-Efficient Revocation in Group Signatures. In *Financial Cryptography 2002*, LNCS 2357, pages 183–197. Springer, 2003.

[34] Nuttapong Attrapadung, Benoît Libert, and Thomas Peters. Computing on authenticated data: New privacy definitions and constructions. In *Advances in Cryptology – ASIACRYPT 2012*, Lecture Notes in Computer Science, pages 367–385. Springer, December 2012. doi:10.1007/978-3-642-34961-4_23.

[35] Man Ho Au, Patrick P. Tsang, Willy Susilo, and Yi Mu. Dynamic Universal Accumulators for DDH Groups and Their Application to Attribute-Based Anonymous Credential Systems. In Marc Fischlin, editor, *CT-RSA*, volume 5473 of *LNCS*, pages 295–308. Springer, 2009. URL: `http://dx.doi.org/10.1007/978-3-642-00862-7_20`, `doi:10.1007/978-3-642-00862-7_20`.

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

[37] Ali Bagherzandi, Stanislaw Jarecki, Nitesh Saxena, and Yanbin Lu. Password-protected secret sharing. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 11: 18th Conference on Computer and Communications Security*, pages 433–444. ACM Press, October 2011.

[38] Josep Balasch, Alfredo Rial, Carmela Troncoso, Bart Preneel, Ingrid Verbauwhede, and Christophe Geuens. Pretp: Privacy-preserving electronic toll pricing. In *USENIX Security Symposium*, pages 63–78. USENIX Association, 2010.

[39] Foteini Baldimtsi and Anna Lysyanskaya. Anonymous credentials light. In *ACM CCS 13: 20th Conference on Computer and Communications Security*, pages 1087–1098. ACM Press, 2013.

[40] Foteini Baldimtsi and Anna Lysyanskaya. On the security of one-witness blind signature schemes. In *Advances in Cryptology – ASIACRYPT 2013, Part II*, volume 8270 of *Lecture Notes in Computer Science*, pages 82–99. Springer, December 2013. `doi:10.1007/978-3-642-42045-0_5`.

[41] Endre Bangerter, Stephan Krenn, Ahmad-Reza Sadeghi, Thomas Schneider, and Joe-Kai Tsay. On the design and implementation of efficient zero-knowledge proofs of knowledge. *Software Performance Enhancements for Encryption and Decryption and Cryptographic Compilers–SPEED-CC*, 9:12–13, 2009.

[42] Michael Barbaro, Tom Zeller, and Saul Hansell. A face is exposed for aol searcher no. 4417749. *New York Times*, 9(2008):8For, 2006.

[43] Niko Baric and Birgit Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In *EUROCRYPT '97*, volume 1233 of LNCS, pages 480–494, 1997.

[44] 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.

[45] Donald Beaver and Stuart Haber. Cryptographic Protocols Provably Secure Against Dynamic Adversaries. In *EUROCRYPT*, pages 307–323, 1992.

[46] Mira Belenkiy, Jan Camenisch, Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Hovav Shacham. Randomizable proofs and delegatable anonymous credentials. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 108–125. Springer, 2009.

[47] Mira Belenkiy, Melissa Chase, Markulf Kohlweiss, and Anna Lysyanskaya. P-signatures and noninteractive anonymous credentials. In Ran Canetti, editor, *TCC 2008: 5th Theory of Cryptography Conference*, volume 4948 of *Lecture Notes in Computer Science*, pages 356–374. Springer, March 2008.

[48] Mira Belenkiy, Melissa Chase, Markulf Kohlweiss, and Anna Lysyanskaya. Compact e-cash and simulatable VRFs revisited. In Hovav Shacham and Brent Waters, editors, *PAIRING 2009: 3rd International Conference on Pairing-based Cryptography*, volume 5671 of *Lecture Notes in Computer Science*, pages 114–131. Springer, August 2009.

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

[50] Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 614–629. Springer, May 2003.

[51] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The One-More-RSA-Inversion Problems and the Security of Chaum's Blind Signature Scheme. *J. Cryptology*, 16(3):185–215, 2003. doi:10.1007/s00145-002-0120-1.

[52] Mihir Bellare and Gregory Neven. Transitive signatures based on factoring and RSA. In *ASIACRYPT '02*, volume 2501 of LNCS, pages 397–414, 2002.

[53] Mihir Bellare and Gregory Neven. Transitive signatures: New schemes and proofs. *IEEE Transactions on Information Theory*, 51:2133–2151, 2005.

[54] Mihir Bellare, Haixia Shi, and Chong Zhang. Foundations of group signatures: The case of dynamic groups. In Alfred Menezes, editor, *Topics in Cryptology – CT-RSA 2005*, volume 3376 of *Lecture Notes in Computer Science*, pages 136–153. Springer, February 2005.

[55] Josh Benaloh and Michael de Mare. One-way Accumulators: A Decentralized Alternative to Digital Signatures. In *EUROCRYPT*, pages 274–285, 1993.

[56] Fabrice Benhamouda, Jan Camenisch, Stephan Krenn, Vadim Lyubashevsky, and Gregory Neven. Better Zero-Knowledge Proofs for Lattice Encryption and Their Application to Group Signatures. In *ASIACRYPT*, 2014.

[57] John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. In *IEEE Symposium on Security and Privacy*, pages 321–334. IEEE Computer Society, 2007.

[58] 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.

[59] Patrik Bichsel, Jan Camenisch, and Franz-Stefan Preiss. A comprehensive framework enabling data-minimizing authentication. In *Proceedings of the 7th ACM workshop on Digital identity management*, pages 13–22. ACM, 2011.

[60] 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.

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

[62] Olivier Blazy, Georg Fuchsbauer, David Pointcheval, and Damien Vergnaud. Signatures on randomizable ciphertexts. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011: 14th* , volume 6571 of *Lecture Notes in Computer Science*, pages 403–422. Springer, March 2011.

[63] Olivier Blazy, David Pointcheval, and Damien Vergnaud. Compact round-optimal partially-blind signatures. In Ivan Visconti and Roberto De Prisco, editors, *SCN 12: 8th International Conference on Security in Communication Networks*, volume 7485 of *Lecture Notes in Computer Science*, pages 95–112. Springer, September 2012.

[64] Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Yvo Desmedt, editor, *PKC 2003: 6th* , volume 2567 of *Lecture Notes in Computer Science*, pages 31–46. Springer, January 2003.

[65] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In *CRYPTO '04*, volume 3152 of LNCS, pages 45–55, 2004.

[66] Dan Boneh and Henry Corrigan-Gibbs. Bivariate Polynomials Modulo Composites and their Applications. In *ASIACRYPT*, 2014. `http://eprint.iacr.org/2014/719`.

[67] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. *SIAM J. Comput.*, 32(3), 2003.

[68] Dan Boneh and David Freeman. Homomorphic signatures for polynomial functions. In *Proc. of Eurocrypt*, 2011. Cryptology ePrint Archive, Report 2011/018.

[69] Dan Boneh and David Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In *Proc. of PKC*, volume 6571 of *LNCS*, pages 1–16, 2011. Cryptology ePrint Archive, Report 2010/453.

[70] Dan Boneh, David Freeman, Jonathan Katz, and Brent Waters. Signing a linear subspace: Signature schemes for network coding. In *Public-Key Cryptography — PKC '09*, volume 5443 of *Springer LNCS*, pages 68–87, 2009.

[71] Dan Boneh and Hovav Shacham. Group Signatures with Verifier-Local Revocation. In *CCS 2004*, pages 168–177. ACM, 2004.

[72] Jason Bordoff and Pascal Noel. Pay-as-you-drive auto insurance: A simple way to reduce driving-related harms and increase equity. *Hamilton Project Discussion Paper*, 2008.

[73] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. Cryptology ePrint Archive, Report 2013/401, 2013. `http://eprint.iacr.org/2013/401`.

[74] John Brainard, Ari Juels, Burton S. Kaliski Jr., and Michael Szydlo. A new two-server approach for authentication with short secrets. In *Proceedings of the 12th USENIX Security Symposium (SECURITY 2003)*, pages 201–214, Washington, DC, USA, August 5–9, 2003. USENIX Association.

[75] Stefan Brands. Untraceable off-line cash in wallets with observers (extended abstract). In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO'93*, volume 773 of *Lecture Notes in Computer Science*, pages 302–318. Springer, August 1994.

[76] Stefan Brands. Restrictive blinding of secret-key certificates. In Louis C. Guillou and Jean-Jacques Quisquater, editors, *Advances in Cryptology – EUROCRYPT'95*, volume 921 of *Lecture Notes in Computer Science*, pages 231–247. Springer, May 1995.

[77] Stefan Brands. *Rethinking Public Key Infrastructure and Digital Certificates— Building in Privacy*. Eindhoven, The Netherlands, 1999.

[78] 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`.

[79] Stefan Brands and Christian Paquin. U-Prove Cryptographic Specification v1. *Microsoft Corporation*, 2010.

[80] Ernest F. Brickell, Jan Camenisch, and Liqun Chen. Direct anonymous attestation. In Vijayalakshmi Atluri, Birgit Pfitzmann, and Patrick McDaniel, editors, *ACM CCS 04: 11th Conference on Computer and Communications Security*, pages 132–145. ACM Press, October 2004.

[81] Ernie Brickell, Jan Camenisch, and Liqun Chen. The DAA scheme in context. In Chris J. Mitchell, editor, *Trusted Computing*, volume 6 of *Professional Applications of Computing*, chapter 5, pages 143–174. Institution of Electrical Engineers, 2005.

[82] Christina Brzuska, Marc Fischlin, Tobias Freudenreich, Anja Lehmann, Marcus Page, Jakob Schelbert, Dominique Schröder, and Florian Volk. Security of sanitizable signatures revisited. In *Public Key Cryptography*, volume 5443 of LNCS, pages 317–336, 2009.

[83] 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`.

[84] Christina Brzuska, Marc Fischlin, Anja Lehmann, and Dominique Schröder. Santizable signatures: How to partially delegate control for authenticated data. In *BIOSIG 2009*, pages 117–128, 2009.

[85] Ahto Buldas, Peeter Laud, and Helger Lipmaa. Accountable Certificate Management Using Undeniable Attestations. In *ACM CCS*, pages 9–17. ACM, 2000.

[86] Ahto Buldas, Peeter Laud, and Helger Lipmaa. Eliminating Counterevidence with Applications to Accountable Certificate Management. *Journal of Computer Security*, 10:2002, 2002.

[87] W. E. Burr, D. F. Dodson, E. M. Newton, R. A. Perlner, W. T. Polk, S. Gupta, and E. A. Nabbus. Electronic authentication guideline. NIST Special Publication 800-63-1, 2011.

[88] Philippe Camacho and Alejandro Hevia. On the impossibility of batch update for cryptographic accumulators. In Michel Abdalla and Paulo S. L. M. Barreto, editors, *Progress in Cryptology - LATINCRYPT 2010: 1st International Conference on Cryptology and Information Security in Latin America*, volume 6212 of *Lecture Notes in Computer Science*, pages 178–188. Springer, August 2010.

[89] Philippe Camacho, Alejandro Hevia, Marcos A. Kiwi, and Roberto Opazo. Strong accumulators from collision-resistant hashing. In Tzong-Chen Wu, Chin-Laung Lei, Vincent Rijmen, and Der-Tsai Lee, editors, *ISC 2008: 11th* , volume 5222 of *Lecture Notes in Computer Science*, pages 471–486. Springer, September 2008.

[90] Jan Camenisch, Maria Dubovitskaya, Robert R. Enderlein, Anja Lehmann, Gregory Neven, Christian Paquin, and Franz-Stefan Preiss. Concepts and languages for privacy-preserving attribute-based authentication. *Journal of Information Security and Applications*, 19(1):25 – 44, 2014. URL: `http://www.sciencedirect.com/science/article/pii/S2214212614000167`, `doi:http://dx.doi.org/10.1016/j.jisa.2014.03.004`.

[91] Jan Camenisch, Maria Dubovitskaya, Robert R. Enderlein, and Gregory Neven. Oblivious transfer with hidden access control from attribute-based encryption. In Ivan Visconti and Roberto De Prisco, editors, *SCN*, volume 7485 of *Lecture Notes in Computer Science*, pages 559–579. Springer, 2012.

[92] Jan Camenisch, Maria Dubovitskaya, Kristiyan Haralambiev, and Markulf Kohlweiss. Unlinkable redactable signatures and their applications to anonymous credentials. In *ASIACRYPT 2015*, 2015.

[93] 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.

[94] Jan Camenisch, Maria Dubovitskaya, and Gregory Neven. Oblivious transfer with access control. In Ehab Al-Shaer, Somesh Jha, and Angelos D. Keromytis, editors, *ACM Conference on Computer and Communications Security*, pages 131–140. ACM, 2009.

[95] Jan Camenisch, Maria Dubovitskaya, Gregory Neven, and Gregory M. Zaverucha. Oblivious transfer with hidden access control policies. In Dario Catalano, Nelly Fazio, Rosario

Gennaro, and Antonio Nicolosi, editors, *PKC 2011: 14th* , volume 6571 of *Lecture Notes in Computer Science*, pages 192–209. Springer, March 2011.

[96] Jan Camenisch, Maria Duvobitskaya, and Alfredo Rial. Uc commitments, revocation, and attribute tokens for privacy-preserving protocol design. Under submission.

[97] Jan Camenisch, Robert R Enderlein, and Gregory Neven. Two-server password-authenticated secret sharing uc-secure against transient corruptions. In *Public-Key Cryptography–PKC 2015*, pages 283–307. Springer, 2015.

[98] 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`.

[99] Jan Camenisch and Thomas Groß. Efficient attributes for anonymous credentials. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *ACM CCS 08: 15th Conference on Computer and Communications Security*, pages 345–356. ACM Press, October 2008.

[100] Jan Camenisch and Thomas Groß. Efficient attributes for anonymous credentials. *ACM Transactions on Information and System Security (TISSEC)*, 15(1):4:1–4:30, 2012.

[101] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. Compact e-cash. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 302–321. Springer, May 2005.

[102] Jan Camenisch, Günter Karjoth, Gregory Neven, and Franz-Stefan Preiss. Anonymously sharing flickr pictures with facebook friends. In *Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society*, pages 13–24. ACM, 2013.

[103] Jan Camenisch, Aggelos Kiayias, and Moti Yung. On the portability of generalized schnorr proofs. In Antoine Joux, editor, *Advances in Cryptology – EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 425–442. Springer, April 2009.

[104] 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`.

[105] 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`.

[106] Jan Camenisch, Maciej Koprowski, and Bogdan Warinschi. Efficient blind signatures without random oracles. In Carlo Blundo and Stelvio Cimato, editors, *SCN 04: 4th International Conference on Security in Communication Networks*, volume 3352 of *Lecture Notes in Computer Science*, pages 134–148. Springer, September 2004.

[107] Jan Camenisch, Stephan Krenn, Anja Lehmann, Gert Læssøe Mikkelsen, Gregory Neven, and Michael Østergaard Pedersen. Formal treatment of privacy-enhancing credential systems. In *SAC 2015*, 2015.

[108] Jan Camenisch, Ioannis Krontiris, Anja Lehmann, Gregory Neven, Christian Paquin, Kai Rannenberg, and Harald Zwingelberg. D2.1 Architecture for Attribute-based Credential Technologies. Technical report, ABC4Trust, 2011.

[109] Jan Camenisch, Ioannis Krontiris, Anja Lehmann, Gregory Neven, Christian Paquin, Kai Rannenberg, and Harald Zwingelberg. H2.1 – abc4trust architecture for developers. ABC4Trust Heartbeat H2.1, 2011. Available from `https://abc4trust.eu`.

[110] Jan Camenisch and Anja Lehmann. (un)linkable pseudonyms for governmental databases. In *ACM CCS 2015*, 2015.

[111] Jan Camenisch, Anja Lehmann, Anna Lysyanskaya, and Gregory Neven. Memento: How to reconstruct your secrets from a single password in a hostile environment. In *Advances in Cryptology–CRYPTO 2014*, pages 256–275. Springer, 2014.

[112] Jan Camenisch, Anja Lehmann, and Gregory Neven. Optimal distributed password verification. In *ACM CCS 2015*, 2015.

[113] Jan Camenisch, Anja Lehmann, Gregory Neven, and Alfredo Rial. Privacy-preserving auditing for attribute-based credentials. In Mirosław Kutyłowski and Jaideep Vaidya, editors, *Computer Security - ESORICS 2014*, volume 8713 of *Lecture Notes in Computer Science*, pages 109–127. Springer International Publishing, 2014. URL: `http://dx.doi.org/10.1007/978-3-319-11212-1_7`, `doi:10.1007/978-3-319-11212-1_7`.

[114] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 93–118. Springer, May 2001.

[115] 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`.

[116] Jan Camenisch and Anna Lysyanskaya. A signature scheme with efficient protocols. In Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano, editors, *SCN 02: 3rd International Conference on Security in Communication Networks*, volume 2576 of *Lecture Notes in Computer Science*, pages 268–289. Springer, September 2002.

[117] Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 56–72. Springer, August 2004.

[118] Jan Camenisch, Anna Lysyanskaya, and Gregory Neven. Practical yet universally composable two-server password-authenticated secret sharing. In *ACM CCS 12: 19th Conference on Computer and Communications Security*, pages 525–536. ACM Press, 2012.

[119] Jan Camenisch, Sebastian Mödersheim, Gregory Neven, Franz-Stefan Preiss, and Alfredo Rial. A prolog program for matching attribute-based credentials to access control policies. Technical Report RZ3890, IBM, 2015.

[120] 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`.

[121] Jan Camenisch, Gregory Neven, and Markus Rückert. Fully anonymous attribute tokens from lattices. In *SCN 2012*, volume 7485 of *Lecture Notes in Computer Science*, pages 57–75. Springer, 2012.

[122] Jan Camenisch, Gregory Neven, and Abhi Shelat. Simulatable adaptive oblivious transfer. In Moni Naor, editor, *Advances in Cryptology – EUROCRYPT 2007*, volume 4515 of *Lecture Notes in Computer Science*, pages 573–590. Springer, May 2007.

[123] Jan Camenisch, Diego A. Ortiz-Yepes, and Franz-Stefan Preiss. Strengthening authentication with privacy-preserving location verification of mobile phones. In *WPES 2015*, 2015.

[124] Jan Camenisch and Els Van Herreweghen. Design and implementation of the idemix anonymous credential system. In Vijayalakshmi Atluri, editor, *ACM CCS 02: 9th Conference on Computer and Communications Security*, pages 21–30. ACM Press, November 2002.

[125] 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.

[126] Sébastien Canard and Amandine Jambert. On extended sanitizable signature schemes. In Josef Pieprzyk, editor, *Topics in Cryptology – CT-RSA 2010*, volume 5985 of *Lecture Notes in Computer Science*, pages 179–194. Springer, March 2010.

[127] Sébastien Canard and Roch Lescuyer. Protecting privacy by sanitizing personal data: a new approach to anonymous credentials. In *ASIACCS 13: 8th Conference on Computer and Communications Security*, pages 381–392. ACM Press, 2013.

[128] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science*, pages 136–145. IEEE Computer Society Press, October 2001.

[129] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th Annual ACM Symposium on Theory of Computing*, pages 209–218. ACM Press, May 1998.

[130] Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In *EUROCRYPT*, pages 255–271, 2003.

[131] Ran Canetti, Shai Halevi, Jonathan Katz, Yehuda Lindell, and Philip D. MacKenzie. Universally composable password-based key exchange. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 404–421. Springer, May 2005.

[132] Dario Catalano and Dario Fiore. Vector commitments and their applications. In *PKC 2013: 16th* , Lecture Notes in Computer Science, pages 55–72. Springer, 2013. `doi:10.1007/978-3-642-36362-7_5`.

[133] Dario Catalano, Dario Fiore, and Mariagrazia Messina. Zero-knowledge sets with short proofs. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 433–450. Springer, April 2008.

[134] Denis Charles, K Jain, and K Lauter. Signatures for network coding. *International Journal of Information and Coding Theory*, 1(1):3–14, 2009.

[135] Melissa Chase. *Efficient non-interactive zero-knowledge proofs for privacy applications*. PhD thesis, Brown University, 2008.

[136] 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/`.

[137] Melissa Chase and Anna Lysyanskaya. On signatures of knowledge. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 78–96. Springer, August 2006.

[138] Melissa Chase, Sarah Meiklejohn, and Greg Zaverucha. Algebraic macs and keyed-verification anonymous credentials. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1205–1216. ACM, 2014.

[139] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.

[140] David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *Advances in Cryptology: Proceedings of CRYPTO '82, Santa Barbara, California, USA, August 23-25, 1982.*, pages 199–203. Plenum Press, New York, 1982.

[141] David Chaum. Blind signature system. In David Chaum, editor, *Advances in Cryptology – CRYPTO'83*, page 153. Plenum Press, New York, USA, 1984.

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

[143] David Chaum, Amos Fiat, and Moni Naor. Untraceable electronic cash. In Shafi Goldwasser, editor, *Advances in Cryptology – CRYPTO'88*, volume 403 of *Lecture Notes in Computer Science*, pages 319–327. Springer, August 1990.

[144] David Chaum and Eugène van Heyst. Group signatures. In *EUROCRYPT*, volume 547 of LNCS, pages 257–265, 1991.

[145] 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`.

[146] George P Chrousos. Stressors, stress, and neuroendocrine integration of the adaptive response: the 1997 hans selye memorial lecture. *Annals of the New York Academy of Sciences*, 851(1):311–335, 1998.

[147] Cheng-Kang Chu and Wen-Guey Tzeng. Efficient k-out-of-n oblivious transfer schemes with adaptive and non-adaptive queries. In Serge Vaudenay, editor, *PKC 2005: 8th* , volume 3386 of *Lecture Notes in Computer Science*, pages 172–183. Springer, January 2005.

[148] 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`.

[149] K. P. L. Coopamootoo and T. Groß. Mental models: An approach to identify privacy concern and behavior. At SOUPS 2014 workshop on Privacy Personas and Segmentation, 2014.

[150] K. P. L. Coopamootoo and T. Groß. Cognitive effort in privacy decision-making vs. 3 x 4: Evaluation of a pilot experiment design. In LASER 2014 Workshop, 2014.

[151] K. P. L. Coopamootoo and T. Groß. Mental models for usable privacy: A position paper. In T.Tryfonas and I.Askoxylakis, editors, *HAS 2014*, volume 8533 of *LNCS*, pages 410–421. Springer Int, 2014.

[152] K. P. L. Coopamootoo and T. Groß. Mental models of online privacy: Structural properties and cognitive maps. In *British HCI 2014*, 2014.

[153] K. P. L. Coopamootoo and T. Groß. Poster: Preliminary investigation of cognitive effort in privacy decision-making: personal information vs. 3 x 4. In *SOUPS 2014*, 2014.

[154] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280 (Proposed Standard), May 2008.

[155] 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.

[156] Scott E. Coull, Matthew Green, and Susan Hohenberger. Controlling access to an oblivious database using stateful anonymous credentials. In Stanislaw Jarecki and Gene Tsudik, editors, *Public Key Cryptography*, volume 5443 of *Lecture Notes in Computer Science*, pages 501–520. Springer, 2009.

[157] Ivan Damgård and Nikos Triandopoulos. Supporting non-membership proofs with bilinear-map accumulators. Cryptology ePrint Archive, Report 2008/538, 2008. http://eprint.iacr.org/2008/538.

[158] Anupam Datta, Ante Derek, John C. Mitchell, and Dusko Pavlovic. Secure protocol composition. In Michael Backes and David A. Basin, editors, *FMSE*, pages 11–23. ACM, 2003.

[159] Hermann de Meer, Manuel Liedel, Henrich C. Pöhls, and Joachim Posegga. Indistinguishability of One-Way Accumulators. Technical Report MIP-1210, Faculty of Computer Science and Mathematics (FIM), University of Passau, 2012.

[160] FILIP De Meyer, Georges De Moor, and L Reed-Fourquet. Privacy protection through pseudonymisation in ehealth. *Studies in health technology and informatics*, 141:111–118, 2007.

[161] Yves-Alexandre de Montjoye, Laura Radaelli, Vivek Kumar Singh, et al. Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science*, 347(6221):536–539, 2015.

[162] 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.

[163] David Derler, Christian Hanser, and Daniel Slamanig. Blank digital signatures: Optimization and practical experiences. In *Privacy and Identity Management for the Future Internet in the Age of Globalisation*, pages 201–215. Springer, 2014.

[164] David Derler, Christian Hanser, and Daniel Slamanig. Privacy-Enhancing Proxy Signatures from Non-interactive Anonymous Credentials. In *Data and Applications Security and Privacy XXVIII - 28th Annual IFIP WG 11.3 Working Conference, DBSec 2014, Vienna, Austria, July 14-16, 2014. Proceedings*, pages 49–65, 2014. URL: http://dx.doi.org/10.1007/978-3-662-43936-4_4, doi:10.1007/978-3-662-43936-4_4.

[165] David Derler, Christian Hanser, and Daniel Slamanig. Revisiting cryptographic accumulators, additional properties and relations to other primitives. In *Topics in Cryptology - CT-RSA 2015, The Cryptographer's Track at the RSA Conference 2015, San Francisco, CA, USA, April 20-24, 2015. Proceedings*, pages 127–144, 2015. URL: http://dx.doi.org/10.1007/978-3-319-16715-2_7, doi:10.1007/978-3-319-16715-2_7.

[166] Mario Di Raimondo and Rosario Gennaro. Provably secure threshold password-authenticated key exchange. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 507–523. Springer, May 2003.

[167] Roger Dingledine, Nick Mathewson, and Paul F. Syverson. Tor: The second-generation onion router. In Matt Blaze, editor, *Proceedings of the 13th USENIX Security Symposium, August 9-13, 2004, San Diego, CA, USA*, pages 303–320. USENIX, 2004.

[168] 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`.

[169] 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`.

[170] Bernice S Elger, Jimison Iavindrasana, Luigi Lo Iacono, Henning Müller, Nicolas Roduit, Paul Summers, and Jessica Wright. Strategies for health data exchange for secondary, cross-institutional clinical research. *Computer methods and programs in biomedicine*, 99(3):230–251, 2010.

[171] EMC Corporation. RSA distributed credential protection. `http://www.emc.com/security/rsa-distributed-credential-protection.htm`.

[172] Alea Fairchild and Bruno de Vuyst. The Evolution of the e-ID card in Belgium: Data Privacy and Multi-Application Usage. In *The Sixth International Conference on Digital Society*, pages 13–16, Valencia, 2012.

[173] Prastudy Fauzi, Helger Lipmaa, and Bingsheng Zhang. Efficient non-interactive zero knowledge arguments for set operations. Cryptology ePrint Archive, Report 2014/006, 2014. `http://eprint.iacr.org/2014/006`.

[174] Nelly Fazio and Antonio Nicolisi. Cryptographic Accumulators: Definitions, Constructions and Applications. Technical report, 2002. URL: `http://cs.nyu.edu/~{}fazio/research/publications/accumulators.pdf`.

[175] Federal Chancellery. The Austrian E-Government Act. *Austrian Federal Law Gazette I*, 7:1–11, 2008. URL: `https://www.ris.bka.gv.at/GeltendeFassung.wxe?Abfrage=Bundesnormen&Gesetzesnummer=20003230`.

[176] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology – CRYPTO'86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, August 1987.

[177] Marc Fischlin. Round-Optimal Composable Blind Signatures in the Common Reference String Model. In *CRYPTO*, volume 4117 of *LNCS*, pages 60–77. Springer, 2006.

[178] Marc Fischlin and Dominique Schröder. On the impossibility of three-move blind signature schemes. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 197–215. Springer, May 2010.

[179] Warwick Ford and Burton S. Kaliski Jr. Server-assisted generation of a strong secret from a password. In *9th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE 2000)*, pages 176–180, Gaithersburg, MD, USA, June 4–16, 2000. IEEE Computer Society.

[180] Cedric Fournet, Markulf Kohlweiss, George Danezis, and Zhengqin Luo. Zql: A compiler for privacy-preserving data processing. In *USENIX Security*, pages 163–178, 2013.

[181] Christina Fragouli and Emina Soljanin. *Network Coding Fundamentals*. Now Publishers Inc., Hanover, MA, USA, 2007.

[182] Georg Fuchsbauer. Commuting signatures and verifiable encryption. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 224–245. Springer, May 2011.

[183] Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig. Practical Round-Optimal Blind Signatures in the Standard Model. In *CRYPTO 2015*, LNCS. Springer, 2015. to appear.

[184] David Galindo and Eric R Verheul. Microdata sharing via pseudonymization. *Work Session on Statistical Data Confidentiality, Manchester*, pages 24–32, 2007.

[185] Sanjam Garg and Divya Gupta. Efficient Round Optimal Blind Signatures. In *EUROCRYPT*, volume 8441 of *LNCS*, pages 477–495. Springer, 2014.

[186] Sanjam Garg, Vanishree Rao, Amit Sahai, Dominique Schröder, and Dominique Unruh. Round optimal blind signatures. In Phillip Rogaway, editor, *Advances in Cryptology – CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 630–648. Springer, August 2011.

[187] Christina Garman, Matthew Green, and Ian Miers. Decentralized anonymous credentials. Cryptology ePrint Archive, Report 2013/622, 2013. `http://eprint.iacr.org/2013/622`.

[188] Rosario Gennaro, Jonathan Katz, Hugo Krawczyk, and Tal Rabin. Secure network coding over the integers. In *Public Key Cryptography — PKC '10*, volume 6056 of *Springer LNCS*, pages 142–160, 2010.

[189] Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. URL: `http://crypto.stanford.edu/craig`.

[190] Essam Ghadafi and Nigel P. Smart. Efficient two-move blind signatures in the common reference string model. In Dieter Gollmann and Felix C. Freiling, editors, *ISC 2012: 15th* , volume 7483 of *Lecture Notes in Computer Science*, pages 274–289. Springer, September 2012.

[191] Esha Ghosh, Olga Ohrimenko, and Roberto Tamassia. Verifiable Member and Order Queries on a List in Zero-Knowledge. Cryptology ePrint Archive, Report 2014/632, 2014. http://eprint.iacr.org/2014/632.

[192] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions (extended abstract). In *FOCS*, pages 464–479, 1984.

[193] Michael T. Goodrich, Roberto Tamassia, and Jasminka Hasic. An Efficient Dynamic and Distributed Cryptographic Accumulator. In *ISC*, volume 2433 of *LNCS*, pages 372–388. Springer, 2002.

[194] Jeremi M. Gosney. Password cracking HPC. Passwords^12 Conference, 2012. URL: http://passwords12.at.ifi.uio.no/Jeremi_Gosney_Password_Cracking_HPC_Passwords12.pdf.

[195] 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.

[196] Thomas Groß. Efficient certification and zero-knowledge proofs of knowledge on infrastructure topology graphs. In *In proceedings of CCSW 2014 ACM Cloud Computing Security Workshop*. ACM Press, nov 2014.

[197] Thomas Groß. Signatures and efficient proofs on committed graphs and NP-statements. In *19th International Conference on Financial Cryptography and Data Security (FC 2015)*, pages 293–314, 2015.

[198] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 415–432. Springer, April 2008.

[199] Vandana Guleria and Ratna Dutta. Issuer-free adaptive oblivious transfer with access policy. In *Information Security and Cryptology-ICISC 2014*, pages 402–418. Springer, 2014.

[200] Vandana Guleria and Ratna Dutta. Universally composable identity based adaptive oblivious transfer with access control. In *Information Security and Cryptology*, pages 109–129. Springer, 2014.

[201] 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.

[202] Joshua D. Guttman. Establishing and preserving protocol security goals. *Journal of Computer Security*, 22(2):203–267, 2014. URL: http://dx.doi.org/10.3233/JCS-140499, doi:10.3233/JCS-140499.

[203] Stuart Haber, Yasuo Hatano, Yoshinori Honda, William Horne, Kunihiko Miyazaki, Tomas Sander, Satoru Tezoku, and Danfeng Yao. Efficient signature schemes supporting redaction, pseudonymization, and data deidentification. In *ASIACCS '08*, pages 353–362, 2008.

[204] Jinguang Han, Willy Susilo, Yi Mu, and Jun Yan. Attribute-based oblivious access control. *The Computer Journal*, 55(10):1202–1215, 2012.

[205] Jinguang Han, Willy Susilo, Yi Mu, and Jun Yan. Efficient oblivious transfers with access control. *Computers & Mathematics with Applications*, 63(4):827–837, 2012.

[206] Christian Hanser and Daniel Slamanig. Blank digital signatures. In *ASIACCS 13: 8th Conference on Computer and Communications Security*, pages 95–106. ACM Press, 2013.

[207] 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.

[208] Christian Hanser and Daniel Slamanig. Structure-Preserving Signatures on Equivalence Classes and their Application to Anonymous Credentials. In *ASIACRYPT*, 2014. Full version: Cryptology ePrint Archive, Report 2014/705.

[209] Carmit Hazay, Jonathan Katz, Chiu-Yuen Koo, and Yehuda Lindell. Concurrently-secure blind signatures without random oracles or setup assumptions. In Salil P. Vadhan, editor, *TCC 2007: 4th Theory of Cryptography Conference*, volume 4392 of *Lecture Notes in Computer Science*, pages 323–341. Springer, February 2007.

[210] James Heather, Gavin Lowe, and Steve Schneider. How to prevent type flaw attacks on security protocols. *Journal of Computer Security*, 11(2):217–244, 2003.

[211] David A Hensher. Electronic toll collection. *Transportation Research Part A: General*, 25(1):9–16, 1991.

[212] Amir Herzberg, Stanislaw Jarecki, Hugo Krawczyk, and Moti Yung. Proactive secret sharing or: How to cope with perpetual leakage. In Don Coppersmith, editor, *Advances in Cryptology – CRYPTO'95*, volume 963 of *Lecture Notes in Computer Science*, pages 339–352. Springer, August 1995.

[213] Alejandro Hevia and Daniele Micciancio. The provable security of graph-based one-time signatures and extensions to algebraic signature schemes. In *ASIACRYPT '02*, volume 2501 of LNCS, pages 379–396, 2002.

[214] Martin Hirt and Kazue Sako. Efficient receipt-free voting based on homomorphic encryption. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 539–556. Springer, May 2000.

[215] Jaap-Henk Hoepman. Revocable privacy. *ENISA Quarterly Review*, 5(2), June 2009.

[216] Jaap-Henk Hoepman and David Galindo. Non-interactive distributed encryption: a new primitive for revocable privacy. In Yan Chen and Jaideep Vaidya, editors, *Proceedings of*

the 10th annual ACM workshop on Privacy in the electronic society, WPES 2011, Chicago, IL, USA, October 17, 2011, pages 81–92. ACM, 2011. doi:10.1145/2046556.2046567.

[217] Home Create Hubs. National strategy for trusted identities in cyberspace. 2010.

[218] Identity Mixer. http://idemix.wordpress.com/.

[219] Malika Izabachène, Benoît Libert, and Damien Vergnaud. Block-wise p-signatures and non-interactive anonymous credentials with efficient attributes. In *Cryptography and Coding*, pages 431–450. Springer, 2011.

[220] David P. Jablon. Password authentication using multiple servers. In David Naccache, editor, *Topics in Cryptology – CT-RSA 2001*, volume 2020 of *Lecture Notes in Computer Science*, pages 344–360. Springer, April 2001.

[221] Stanislaw Jarecki, Aggelos Kiayias, and Hugo Krawczyk. Round-optimal password-protected secret sharing and t-pake in the password-only model. In *Advances in Cryptology–ASIACRYPT 2014*, pages 233–253. Springer, 2014.

[222] Stanislaw Jarecki and Xiaomin Liu. Fast secure computation of set intersection. In Juan A. Garay and Roberto De Prisco, editors, *SCN 10: 7th International Conference on Security in Communication Networks*, volume 6280 of *Lecture Notes in Computer Science*, pages 418–435. Springer, September 2010.

[223] Oliver P John and Sanjay Srivastava. The big five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research*, 2(1999):102–138, 1999.

[224] Robert Johnson, David Molnar, Dawn Song, and David Wagner. Homomorphic signature schemes. In *CT-RSA*, pages 244–262. Springer-Verlag, 2002.

[225] Daniel Kahneman. *Attention and effort*. Citeseer, 1973.

[226] Daniel Kahneman. *Thinking fast and slow*. Farrar, Strauss, 2011.

[227] Daniel Kahneman and Jackson Beatty. Pupil diameter and load on memory. *Science*, 154(3756), 1966.

[228] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *Advances in Cryptology – ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 177–194. Springer, December 2010.

[229] Jonathan Katz, Philip D. MacKenzie, Gelareh Taban, and Virgil D. Gligor. Two-server password-only authenticated key exchange. In John Ioannidis, Angelos Keromytis, and Moti Yung, editors, *ACNS 05: 3rd International Conference on Applied Cryptography and Network Security*, volume 3531 of *Lecture Notes in Computer Science*, pages 1–16. Springer, June 2005.

[230] Aggelos Kiayias and Moti Yung. Group signatures with efficient concurrent join. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 198–214. Springer, May 2005.

[231] Aggelos Kiayias and Hong-Sheng Zhou. Concurrent blind signatures without random oracles. In Roberto De Prisco and Moti Yung, editors, *SCN 06: 5th International Conference on Security in Communication Networks*, volume 4116 of *Lecture Notes in Computer Science*, pages 49–62. Springer, September 2006.

[232] Markulf Kohlweiss, Sebastian Faust, Lothar Fritsch, Bartek Gedrojc, and Bart Preneel. Efficient oblivious augmented maps: Location-based services with a payment broker. In Nikita Borisov and Philippe Golle, editors, *Privacy Enhancing Technologies*, volume 4776 of *Lecture Notes in Computer Science*, pages 77–94. Springer, 2007.

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

[234] M. Krohn, M. Freedman, and D. Mazieres. On-the-fly verification of rateless erasure codes for efficient content distribution. In *Proc. of IEEE Symposium on Security and Privacy*, pages 226–240, 2004.

[235] Jorn Lapon, Markulf Kohlweiss, Bart De Decker, and Vincent Naessens. Analysis of revocation strategies for anonymous idemix credentials. In *CMS*, pages 3–17. Springer, 2011.

[236] Anja Lehmann, Patrik Bichsel, Patrik Bichsel, Bud Bruegger, Jan Camenisch, Alberto Crespo Garcia, Thomas Gross, André Gutwirth, Moritz Horsch, Detlef Houdeau, Detlef Hühnlein, Frank-Michael Kamm, Stephan Krenn, Gregory Neven, Charles Bastos Rodriguez, Johannes Schmölz, and Charlotte Bolliger. Survey and Analysis of Existing eID and Credential Systems. Technical Report Deliverable D32.1, FutureID, 2013.

[237] Herbert Leitold, Arno Hollosi, and Reinhard Posch. Security Architecture of the Austrian Citizen Card Concept. In *ACSAC 2002*, pages 391–402, 2002.

[238] Matt Lepinksi, Silvio Micali, et al. Collusion-free protocols. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 543–552. ACM, 2005.

[239] Allison B. Lewko, Tatsuaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In *EUROCRYPT*, 2010.

[240] Jiangtao Li, Ninghui Li, and Rui Xue. Universal accumulators with efficient nonmembership proofs. In Jonathan Katz and Moti Yung, editors, *ACNS 07: 5th International Conference on Applied Cryptography and Network Security*, volume 4521 of *Lecture Notes in Computer Science*, pages 253–269. Springer, June 2007.

[241] Jin Li, Man Ho Au, Willy Susilo, Dongqing Xie, and Kui Ren. Attribute-based signature and its applications. In Dengguo Feng, David A. Basin, and Peng Liu, editors, *ASIACCS 10: 5th Conference on Computer and Communications Security*, pages 60–69. ACM Press, April 2010.

[242] 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.

[243] Benoît Libert and Moti Yung. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In Daniele Micciancio, editor, *TCC 2010: 7th Theory of Cryptography Conference*, volume 5978 of *Lecture Notes in Computer Science*, pages 499–517. Springer, February 2010.

[244] Yehuda Lindell. Bounded-concurrent secure two-party computation without setup assumptions. In *35th Annual ACM Symposium on Theory of Computing*, pages 683–692. ACM Press, June 2003.

[245] Helger Lipmaa. Secure accumulators from euclidean rings without trusted setup. In Feng Bao, Pierangela Samarati, and Jianying Zhou, editors, *ACNS 12: 10th International Conference on Applied Cryptography and Network Security*, volume 7341 of *Lecture Notes in Computer Science*, pages 224–240. Springer, June 2012.

[246] Wouter Lueks, Gergely Alpár, Jaap-Henk Hoepman, and Pim Vullers. Fast revocation of attribute-based credentials for both users and verifiers. In *ICT Systems Security and Privacy Protection*, pages 463–478. Springer, 2015.

[247] Wouter Lueks, Maarten H. Everts, and Jaap-Henk Hoepman. Revocable privacy: Principles, use cases, and technologies. Annual Privacy Forum (APF 2015) (to appear).

[248] Wouter Lueks, Maarten H. Everts, and Jaap-Henk Hoepman. Revocable privacy 2012 – use cases. Technical Report 35627, TNO, 2012.

[249] Wouter Lueks, Jaap-Henk Hoepman, and Klaus Kursawe. Forward-secure distributed encryption. In Emiliano De Cristofaro and Steven J. Murdoch, editors, *Privacy Enhancing Technologies - 14th International Symposium, PETS 2014, Amsterdam, The Netherlands, July 16-18, 2014. Proceedings*, volume 8555 of *Lecture Notes in Computer Science*, pages 123–142. Springer, 2014. doi:10.1007/978-3-319-08506-7_7.

[250] Anna Lysyanskaya, Ronald L. Rivest, Amit Sahai, and Stefan Wolf. Pseudonym systems. In Howard M. Heys and Carlisle M. Adams, editors, *SAC 1999: 6th Annual International Workshop on Selected Areas in Cryptography*, volume 1758 of *Lecture Notes in Computer Science*, pages 184–199. Springer, August 2000.

[251] 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.

[252] Philip MacKenzie and Ke Yang. On simulation-sound trapdoor commitments. In *EUROCRYPT 2004*, pages 382–400. Springer, 2004.

[253] Philip D. MacKenzie, Thomas Shrimpton, and Markus Jakobsson. Threshold password-authenticated key exchange. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 385–400. Springer, August 2002.

[254] Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. Attribute-based signatures. In Aggelos Kiayias, editor, *Topics in Cryptology – CT-RSA 2011*, volume 6558 of *Lecture Notes in Computer Science*, pages 376–392. Springer, February 2011.

[255] Atefeh Mashatan and Serge Vaudenay. A Fully Dynamic Universal Accumulator. *Proceedings of the Romanian Academy*, 14:269–285, 2013.

[256] S Massoud Amin and Bruce F Wollenberg. Toward a smart grid: power delivery for the 21st century. *Power and Energy Magazine, IEEE*, 3(5):34–41, 2005.

[257] Sarah Meiklejohn, Keaton Mowery, Stephen Checkoway, and Hovav Shacham. The phantom tollbooth: Privacy-preserving electronic toll collection in the presence of driver collusion. In *USENIX Security Symposium*, volume 201, 2011.

[258] Sarah Meiklejohn, Hovav Shacham, and David Mandell Freeman. Limitations on Transformations from Composite-Order to Prime-Order Groups: The Case of Round-Optimal Blind Signatures. In Abe [11], pages 519–538.

[259] Silvio Micali, Michael O. Rabin, and Joe Kilian. Zero-knowledge sets. In *44th Annual Symposium on Foundations of Computer Science*, pages 80–91. IEEE Computer Society Press, October 2003.

[260] Silvio Micali and Ronald L. Rivest. Transitive signature schemes. In Bart Preneel, editor, *Topics in Cryptology – CT-RSA 2002*, volume 2271 of *Lecture Notes in Computer Science*, pages 236–243. Springer, February 2002.

[261] Ian Miers, Christina Garman, Matthew Green, and Aviel D. Rubin. Zerocoin: Anonymous distributed E-cash from bitcoin. In *2013 IEEE Symposium on Security and Privacy*, pages 397–411. IEEE Computer Society Press, 2013.

[262] Kunihiko Miyazaki, Goichiro Hanaoka, and Hideki Imai. Digitally signed document sanitizing scheme based on bilinear maps. In *ASIACCS '06: Proceedings of the 2006 ACM Symposium on Information, computer and communications security*, pages 343–354, 2006.

[263] S. Mödersheim, O. Almousa, Bud Bruegger, and Max Tuengerthal. A Formal Verification of the FutureID architecture. Technical report, DTU Compute-2015-, 2015. Available at `https://dms-prext.fraunhofer.de/livelink/livelink.exe/overview/6419289`.

[264] Sebastian Mödersheim. Deciding Security for a Fragment of ASLan. In *ESORICS*, pages 127–144. Springer, 2012. URL: `http://dx.doi.org/10.1007/978-3-642-33167-1_8`, `doi:10.1007/978-3-642-33167-1_8`.

[265] Sebastian Mödersheim, Thomas Groß, and Luca Viganò. Defining privacy is supposed to be easy. In *Logic for Programming, Artificial Intelligence, and Reasoning - 19th International Conference, LPAR-19, Stellenbosch, South Africa, December 14-19, 2013. Proceedings*, pages 619–635, 2013.

[266] Sebastian Modersheim and Georgios Katsoris. A sound abstraction of the parsing problem. In *Computer Security Foundations Symposium (CSF), 2014 IEEE 27th*, pages 259–273. IEEE, 2014.

[267] Sebastian Mödersheim and Luca Viganò. Sufficient conditions for vertical composition of security protocols. In *Proceedings of the 9th ACM symposium on Information, computer and communications security*, pages 435–446. ACM, 2014.

[268] 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`.

[269] P. Modesti. AnBx: Automatic generation and verification of security protocols implementations. In *8th International Symposium on Foundations & Practice of Security*. to appear, 2015.

[270] Paolo Modesti. Efficient Java code generation of security protocols specified in AnB/AnBx. In *Security and Trust Management*, pages 204–208. Springer International Publishing, 2014.

[271] Collin Mulliner, Ravishankar Borgaonkar, Patrick Stewin, and Jean Pierre Seifert. SMS-based one-time passwords: Attacks and defense (short paper). In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 7967 LNCS, pages 150–159, 2013. `doi:10.1007/978-3-642-39235-1\_9`.

[272] 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`.

[273] Moni Naor and Benny Pinkas. Oblivious transfer with adaptive queries. In *CRYPTO*, pages 573–590, 1999.

[274] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy*, pages 111–125. IEEE Computer Society Press, May 2008.

[275] Ingo Naumann and Giles Hogben. Privacy features of European eID card specifications. *Network Security*, 2008(8):9–13, 2008. `doi:http://dx.doi.org/10.1016/S1353-4858(08)70097-7`.

[276] Thomas Neubauer and Johannes Heurix. A methodology for the pseudonymization of medical data. *International journal of medical informatics*, 80(3):190–204, 2011.

[277] Gregory Neven. A simple transitive signature scheme for directed trees. *Theoretical Computer Science*, 396(1-3):277–282, 2008.

[278] Lan Nguyen. Accumulators from bilinear pairings and applications. In Alfred Menezes, editor, *Topics in Cryptology – CT-RSA 2005*, volume 3376 of *Lecture Notes in Computer Science*, pages 275–292. Springer, February 2005.

[279] Lan Nguyen and Christian Paquin. U-prove designated-verifier accumulator revocation extension. Technical report, Tech. Rep. MSR-TR-2014-85, Microsoft Research, 2014.

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

[281] Takashi Nishide, Kazuki Yoneyama, and Kazuo Ohta. Attribute-based encryption with partially hidden encryptor-specified access structures. In Steven M. Bellovin, Rosario Gennaro, Angelos D. Keromytis, and Moti Yung, editors, *ACNS*, volume 5037 of *Lecture Notes in Computer Science*, pages 111–129, 2008.

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

[283] Kaisa Nyberg. Commutativity in Cryptography. In *1st International Trier Conference in Functional Analysis*. Walter Gruyter & Co, 1996.

[284] Kaisa Nyberg. Fast accumulated hashing. In Dieter Gollmann, editor, *Fast Software Encryption – FSE'96*, volume 1039 of *Lecture Notes in Computer Science*, pages 83–87. Springer, February 1996.

[285] OECD. National Strategies and Policies for Digital Identity Management in OECD Countries, 2011.

[286] KW Ogden. Privacy issues in electronic toll collection. *Transportation Research Part C: Emerging Technologies*, 9(2):123–134, 2001.

[287] Tatsuaki Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In *Advances in Cryptology – CRYPTO '92*, volume 740, pages 31–53, 1992.

[288] Tatsuaki Okamoto. Efficient blind and partially blind signatures without random oracles. In Shai Halevi and Tal Rabin, editors, *TCC 2006: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 80–99. Springer, March 2006.

[289] Oracle. Java™ Cryptography Architecture (JCA) Reference Guide. `http://docs.oracle.com/javase/7/docs/technotes/guides/security/crypto/CryptoSpec.html`.

[290] Clemens Orthacker, Martin Centner, and Christian Kittl. Qualified Mobile Server Signature. In *Proceedings of the 25th TC 11 International Information Security Conference SEC 2010*, 2010.

[291] Diego Alejandro Ortiz-Yepes. Enhancing authentication in ebanking with nfc-enabled mobile phones. *ERCIM News*, 76:63–64, 2009.

[292] Raphael Overbeck. A Step Towards QC Blind Signatures. Cryptology ePrint Archive, Report 2009/102, 2009. `http://eprint.iacr.org/`.

[293] Christian Paquin and Greg Zaverucha. U-prove cryptographic specification v1. 1. Technical report, Microsoft Technical Report, http://connect. microsoft. com/site1188, 2011.

[294] Kun Peng and Feng Bao. Vulnerability of a non-membership proof scheme. In *SECRYPT*, pages 1–4, July 2010.

[295] Birgit Pfitzmann, editor. *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*. Springer, 2001.

[296] Henrich C. Pöhls, Stefan Peters, Kai Samelin, Joachim Posegga, and Hermann de Meer. Malleable Signatures for Resource Constrained Platforms. In *WISTP*, volume 7886 of *LNCS*, pages 18–33. Springer, 2013.

[297] Henrich C. Pöhls and Kai Samelin. On Updatable Redactable Signatures. In *ACNS*, volume 8479 of *LNCS*. Springer, 2014.

[298] K Pommerening, M Reng, P Debold, and S Semler. Pseudonymization in medical research-the generic data protection concept of the tmf. *GMS Medizinische Informatik, Biometrie und Epidemiologie*, 1(3):2005–1, 2005.

[299] Raluca A Popa, Hari Balakrishnan, and Andrew J Blumberg. Vpriv: Protecting privacy in location-based vehicular services. In *USENIX security symposium*, pages 335–350, 2009.

[300] Niels Provos and David Mazières. A future-adaptable password scheme. In *USENIX Annual Technical Conference, FREENIX Track*, pages 81–91. USENIX, 1999.

[301] Michael O. Rabin. How to exchange secrets by oblivious transfer. 1981.

[302] Christof Rath, Simon Roth, Harald Bratko, and Thomas Zefferer. Encryption-based second authentication factor solutions for qualified server-side signature creation. In *4th International Conference on Electronic Government and the Information Systems Perspective, EGOVIS 2015*, 2015. in press.

[303] Christof Rath, Simon Roth, Manuel Schallar, and Thomas Zefferer. A secure and flexible server-based mobile eID and e-signature solution. In *Proceedings of the 8th International Conference on Digital Society, ICDS 2014, Barcelona, Spain*, pages 7 – 12. IARIA, 2014.

[304] Alfredo Rial. Blind attribute-based encryption and oblivious transfer with fine-grained access control. *Designs, Codes and Cryptography*.

[305] Alfredo Rial and George Danezis. Privacy-preserving smart metering. In Yan Chen and Jaideep Vaidya, editors, *WPES*, pages 49–60. ACM, 2011.

[306] Alfredo Rial, George Danezis, and Markulf Kohlweiss. Privacy-preserving smart metering revisited.

[307] Alfredo Rial and Bart Preneel. Blind attribute-based encryption and oblivious transfer with fine-grained access control. *COSIC technical report*, 2010.

[308] Ronald Rivest. Two signature schemes. Slides from talk given at Cambridge University, 2000. `http://people.csail.mit.edu/rivest/Rivest-CambridgeTalk.pdf`.

[309] Ronald L. Rivest, Adi Shamir, and Yael Tauman. How to leak a secret: Theory and applications of ring signatures. In *Essays in Memory of Shimon Even*, pages 164–186, 2006.

[310] Markus Rückert. Lattice-based blind signatures. In Masayuki Abe, editor, *Advances in Cryptology – ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 413–430. Springer, December 2010.

[311] Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *Advances in Cryptology–EUROCRYPT 2005*, pages 457–473. Springer, 2005.

[312] 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`.

[313] Tomas Sander. Efficient accumulators without trapdoor extended abstracts. In Vijay Varadharajan and Yi Mu, editors, *ICICS 99: 2nd International Conference on Information and Communication Security*, volume 1726 of *Lecture Notes in Computer Science*, pages 252–262. Springer, November 1999.

[314] Tomas Sander, Amnon Ta-Shma, and Moti Yung. Blind, auditable membership proofs. In Yair Frankel, editor, *FC 2000: 4th* , volume 1962 of *Lecture Notes in Computer Science*, pages 53–71. Springer, February 2000.

[315] Bruce Schneier. What Our Top Spy Doesn't Get: Security and Privacy Aren't Opposites. *Wired*, January 2008.

[316] Jae Hong Seo and Jung Hee Cheon. Beyond the Limitation of Prime-Order Bilinear Groups, and Round Optimal Blind Signatures. In *TCC*, volume 7194 of *LNCS*, pages 133–150. Springer, 2012.

[317] Siamak Fayyaz Shahandashti and Reihaneh Safavi-Naini. Threshold attribute-based signatures and their application to anonymous credential systems. In Bart Preneel, editor, *AFRICACRYPT 09: 2nd International Conference on Cryptology in Africa*, volume 5580 of *Lecture Notes in Computer Science*, pages 198–216. Springer, June 2009.

[318] Siamak Fayyaz Shahandashti, Mahmoud Salmasizadeh, and Javad Mohajeri. A provably secure short transitive signature scheme from bilinear group pairs. In *Security and Communication Networks*, volume 3352 of *LNCS*, pages 60–76, 2005.

[319] Daniel Slamanig. Dynamic accumulator based discretionary access control for outsourced storage with unlinkable access - (short paper). In Angelos D. Keromytis, editor, *FC 2012: 16th* , volume 7397 of *Lecture Notes in Computer Science*, pages 215–222. Springer, February / March 2012.

[320] Daniel Slamanig, Raphael Spreitzer, and Thomas Unterluggauer. Adding Controllable Linkability to Pairing-Based Group Signatures For Free. In *17th International Conference on Information Security, ISC 2014*, 2014. in press.

[321] Daniel Slamanig, Klaus Stranacher, and Bernd Zwattendorfer. User-centric identity as a service-architecture for eids with selective attribute disclosure. In *19th ACM Symposium on Access Control Models and Technologies, SACMAT '14, London, ON, Canada - June 25 - 27, 2014*, pages 153–164, 2014. URL: http://doi.acm.org/10.1145/2613087.2613093, doi:10.1145/2613087.2613093.

[322] N. P. Smart and F. Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. In *Public Key Cryptography — PKC '10*, volume 6056 of *Springer LNCS*, pages 420–443, 2010.

[323] Marcus Stadler. *Cryptographic Protocols for Revocable Privacy*. PhD thesis, Swiss Federal Institute of Technology, Zürich, 1996.

[324] Ron Steinfeld, Laurence Bull, and Yuliang Zheng. Context extraction signatures. In *Information Security and Cryptology (ICISC)*, volume 2288 of LNCS, pages 285–304, 2001.

[325] Amang Sudarsono, Toru Nakanishi, and Nobuo Funabiki. Efficient Proofs of Attributes in Pairing-Based Anonymous Credential System. In Simone Fischer-Hübner and Nicholas Hopper, editors, *PETS*, volume 6794 of *Lecture Notes in Computer Science*, pages 246–263. Springer, 2011.

[326] 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.

[327] Michael Szydlo and Burton S. Kaliski Jr. Proofs for two-server password authentication. In Alfred Menezes, editor, *Topics in Cryptology – CT-RSA 2005*, volume 3376 of *Lecture Notes in Computer Science*, pages 227–244. Springer, February 2005.

[328] The European Parliament and the Council of the European Union. DIRECTIVE 1999/93/EC OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 13 December 1999 on a Community framework for electronic signatures, 1999. URL: http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2000:013:0012:0020:EN:PDF.

[329] 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.

[330] Carmela Troncoso, George Danezis, Eleni Kosta, and Bart Preneel. Pripayd: privacy friendly pay-as-you-drive insurance. In Peng Ning and Ting Yu, editors, *WPES*, pages 99–107. ACM, 2007.

[331] Patrick P. Tsang, Man Ho Au, Apu Kapadia, and Sean W. Smith. Blacklistable anonymous credentials: blocking misbehaving users without ttps. In Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson, editors, *Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007, Alexandria, Virginia, USA, October 28-31, 2007*, pages 72–81. ACM, 2007. doi:10.1145/1315245.1315256.

[332] Patrick P. Tsang, Apu Kapadia, Cory Cornelius, and Sean W. Smith. Nymble: Blocking misbehaving users in anonymizing networks. *IEEE Trans. Dependable Sec. Comput.*, 8(2):256–269, 2011. doi:10.1109/TDSC.2009.38.

[333] Gene Tsudik and Shouhuai Xu. Accumulating composites and improved group signing. In Chi-Sung Laih, editor, *Advances in Cryptology – ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, pages 269–286. Springer, November / December 2003.

[334] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In *Advances in Cryptology — EUROCRYPT '10*, volume 6110 of *Springer LNCS*, pages 24–43, 2010.

[335] Eric R. Verheul. Self-blindable credential certificates from the Weil pairing. In Colin Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 533–551. Springer, December 2001.

[336] Pim Vullers and Gergely Alpár. Efficient Selective Disclosure on Smart Cards Using Idemix. In *IDMAN 2013*, IFIP AICT 396, pages 53–67, 2013.

[337] Peishun Wang, Huaxiong Wang, and Josef Pieprzyk. A new dynamic accumulator for batch updates. In Sihan Qing, Hideki Imai, and Guilin Wang, editors, *ICICS 07: 9th International Conference on Information and Communication Security*, volume 4861 of *Lecture Notes in Computer Science*, pages 98–112. Springer, December 2007.

[338] Brent Waters. Ciphertext-policy attribute-based encryption: an expressive, efficient, and provably secure realization. In *Public Key Cryptography — PKC '11*, pages 53–70, 2011.

[339] David Watson and Lee Anna Clark. The panas-x: Manual for the positive and negative affect schedule – expanded form. Technical report, University of Iowa, Department of Psychology, 1999.

[340] Lei Wei, Scott E. Coull, and Michael K. Reiter. Bounded vector signatures and their applications. In *ASIACCS '11*, pages 277–285, 2011.

[341] 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.

[342] Ling-Ling Xu and Fang-Guo Zhang. Oblivious transfer with threshold access control. *Journal of Information Science and Engineering*, 28(3):555–570, 2012.

[343] Lingling Xu and Fangguo Zhang. Oblivious transfer with complex attribute-based access control. In *Information Security and Cryptology-ICISC 2010*, pages 370–395. Springer, 2011.

[344] Lingling Xu, Fangguo Zhang, and Yamin Wen. Oblivious transfer with access control and identity-based encryption with anonymous key issuing. *Journal of Electronics (China)*, 28(4-6):571–579, 2011.

[345] Xun Yi. Directed transitive signature scheme. In *CT-RSA '07*, volume 4377 of LNCS, pages 129–144, 2007.

[346] Greg Zaverucha. U-prove id escrow extension. Technical report, TechReport MSR-TR-2013-86, 2013.

[347] Thomas Zefferer and Vesna Krnjic. Usability Evaluation of Electronic Signature Based E-Government Solutions. In *Proceedings of the IADIS International Conference WWW/INTERNET 2012*, pages 227–234, 2012.

[348] Ye Zhang, Man Ho Au, Duncan S. Wong, Qiong Huang, Nikos Mamoulis, David W. Cheung, and Siu-Ming Yiu. Oblivious transfer with access control : Realizing disjunction without duplication. In Marc Joye, Atsuko Miyaji, and Akira Otsuka, editors, *Pairing*, volume 6487 of *Lecture Notes in Computer Science*, pages 96–115. Springer, 2010.

[349] Fang Zhao, Ton Kalker, Muriel Médard, and Keesook Han. Signatures for content distribution with network coding. In *Proc. Intl. Symp. Info. Theory (ISIT)*, 2007.

[350] 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.

[351] 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.

[352] Bernd Zwattendorfer and Daniel Slamanig. Design strategies for a privacy-friendly austrian eid system in the public cloud. *Computers & Security*, (0):–, 2015. URL: `http://www.sciencedirect.com/science/article/pii/S0167404815000346`, `doi:http://dx.doi.org/10.1016/j.cose.2015.03.002`.